

EXHIBIT A

**THIS EXHIBIT HAS BEEN
REDACTED IN ITS ENTIRETY**

EXHIBIT B

Live Traffic Analysis of TCP/IP Gateways

Page 1 of 22

Internet Society's *Networks and Distributed Systems Security Symposium*, March 1998.

Live Traffic Analysis of TCP/IP Gateways

Phillip A. Porras
porras@csdl.sri.com
Computer Science Laboratory

SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

Alfonso Valdes
avaldes@csdl.sri.com
Electromagnetic and Remote
Sensing Laboratory
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

*The work presented in this paper is currently funded by
DARPA/ITO under contract number F30602-96-C-0294.*

Point of Contact: Phillip A. Porras
Phone: (415) 859-3232
Fax: (415) 859-2844

November 10 1997

ABSTRACT

We enumerate a variety of ways to extend both statistical and signature-based intrusion-detection analysis techniques to monitor network traffic. Specifically, we present techniques to analyze TCP/IP packet streams that flow through network gateways for signs of malicious activity, nonmalicious failures,

<http://www.sdl.sri.com/projects/emerald/live-traffic.html>

10/22/2004

SYM_P_0068844

and other exceptional events. The intent is to demonstrate, by example, the utility of introducing gateway surveillance mechanisms to monitor network traffic. We present this discussion of gateway surveillance mechanisms as complementary to the filtering mechanisms of a large enterprise network, and illustrate the usefulness of surveillance in directly enhancing the security and stability of network operations.

1. Introduction

Mechanisms for parsing and filtering hostile external network traffic [2],[4] that could reach internal network services have become widely accepted as prerequisites for limiting the exposure of internal network assets while maintaining interconnectivity with external networks. The encoding of filtering rules for packet- or transport-layer communication should be enforced at entry points between internal networks and external traffic. Developing filtering rules that strike an optimal balance between the restrictiveness necessary to suppress the entry of unwanted traffic, while allowing the necessary flows demanded for user functionality, can be a nontrivial exercise [3].

In addition to intelligent filtering, there have been various developments in recent years in passive surveillance mechanisms to monitor network traffic for signs of malicious or anomalous (e.g., potentially erroneous) activity. Such tools attempt to provide network administrators timely insight into noteworthy exceptional activity. Real-time monitoring promises an added dimension of control and insight into the flow of traffic between the internal network and its external environment. The insight gained through fielded network traffic monitors could also aid sites in enhancing the effectiveness of their firewall filtering rules.

However, traffic monitoring is not a free activity—especially live traffic monitoring. In presenting our discussion of network analysis techniques, we fully realize the costs they imply with respect to computational resources and human oversight. For example, obtaining the necessary input for surveillance involves the deployment of instrumentation to parse, filter, and format event streams derived from potentially high-volume packet transmissions. Complex event analysis, response logic, and human management of the analysis units also introduce costs. Clearly, the introduction of network surveillance

mechanisms on top of already-deployed protective traffic filters is an expense that requires justification. In this paper, we outline the benefits of our techniques and seek to persuade the reader that the costs can be worthwhile.

2. Toward Generalized Network Surveillance

The techniques presented in this paper are extensions of earlier work by SRI in developing analytical methods for detecting anomalous or known intrusive activity [1], [5], [12], [13]. Our earlier intrusion-detection efforts in developing IDDES (Intrusion Detection Expert System) and later NIDES (Next-Generation Intrusion Detection Expert System) were oriented toward the surveillance of user-session and host-layer activity. This previous focus on session activity within host boundaries is understandable given that the primary input to intrusion-detection tools, audit data, is produced by mechanisms that tend to be locally administered within a single host or domain. However, as the importance of network security has grown, so too has the need to expand intrusion-detection technology to address network infrastructure and services. In our current research effort, EMERALD (Event Monitoring Enabling Responses to Anomalous Live Disturbances), we explore the extension of our intrusion-detection methods to the analysis of network activity.

Network monitoring, in the context of fault detection and diagnosis for computer network and telecommunication environments, has been studied extensively by the network management and alarm correlation community [8], [11], [15], [16]. The high-volume distributed event correlation technology promoted in some projects provides an excellent foundation for building truly scalable network-aware surveillance technology for misuse. However, these efforts focus primarily on the health and status (fault detection and/or diagnosis) or performance of the target network, and do not cover the detection of intentionally abusive traffic. Indeed, some simplifications in the fault analysis and diagnosis community (e.g., assumptions of stateless correlation, which precludes event ordering; simplistic time-out metrics for resetting the tracking of problems; ignoring individuals/sources responsible for exceptional activity) do not translate well to a malicious environment for detecting intrusions.

Earlier work in the intrusion-detection community attempting to address the issue of network surveillance includes the Network Security Monitor (NSM), developed at UC Davis [6], and the Network Anomaly Detection and Intrusion Reporter (NADIR) [7], developed at Los Alamos National Laboratory (LANL). Both performed broadcast LAN packet monitoring to analyze traffic patterns for known hostile or anomalous activity.[1] Further research by UC Davis in the Distributed Intrusion Detection System (DIDS) [23] and later Graph-based Intrusion Detection System (GRIDS) [25] projects has attempted to extend intrusion monitoring capabilities beyond LAN analysis, to provide multi-LAN and very large-scale network coverage.

This paper takes a pragmatic look at the issue of packet and/or datagram analysis based on statistical anomaly detection and signature-analysis techniques. This work is being performed in the context of SRI's latest intrusion-detection effort, EMERALD, a distributed scalable tool suite for tracking malicious activity through and across large networks [20].

EMERALD introduces a building-block approach to network surveillance, attack isolation, and automated response. The approach employs highly distributed, independently tunable, surveillance and response monitors that are deployable polymorphically at various abstract layers in a large network. These monitors demonstrate a streamlined intrusion-detection design that combines signature analysis with statistical profiling to provide localized real-time protection of the most widely used network services and components on the Internet.

Among the general types of analysis targets that EMERALD monitors are network gateways. We describe several analysis techniques that EMERALD implements, and discuss their use in analyzing malicious, faulty, and other exceptional network activity. EMERALD's surveillance modules will monitor entry points that separate external network traffic from an enterprise network and its constituent local domains.[3] We present these surveillance techniques as complementary to the filtering mechanisms of a large enterprise network, and illustrate their utility in directly enhancing the security and stability of network operations.

We first consider the candidate event streams that pass through network entry points. Critical to the effective monitoring of operations is the careful selection and organization of these event streams such that an analysis based on a selected event stream will provide meaningful insight into the target activity. We identify effective analytical techniques for processing the event stream given specific analysis objectives. Sections 4 and 5 explore how both statistical anomaly detection and signature analysis can be applied to identify activity worthy of review and possible response. All such claims are supported by examples. More broadly, in Section 6 we discuss the correlation of analysis results produced by surveillance components deployed independently throughout the entry points of our protected intranet. We discuss how events of limited significance to a local surveillance monitor may be aggregated with results from other strategically deployed monitors to provide insight into more wide-scale problems or threats against the intranet. Section 7 discusses the issue of response.

3. Event Stream Selection

The success or failure of event analysis should be quantitatively measured for qualities such as accuracy and performance; both are assessable through testing. A more difficult but equally important metric to assess is completeness. With regard to network surveillance, inaccuracy is reflected in the number of legitimate transactions flagged as abnormal or malicious (false positives), incompleteness is reflected in the number of harmful transactions that escape detection (false negatives), and performance is measured by the rate at which transactions can be processed. All three measurements of success or failure directly depend on the quality of the event stream upon which the analysis is based. Here, we consider the objective of providing real-time surveillance of TCP/IP-based networks for malicious or exceptional network traffic. In particular, our network surveillance mechanisms can be integrated onto, or interconnected with, network gateways that filter traffic between a protected intranet and external networks.

IP traffic represents an interesting candidate event stream for analysis. Individually, packets represent parsable activity records, where key data within the header and data segment can be statistically analyzed and/or heuristically parsed for response-worthy activity. However, the sheer volume of potential packets dictates careful assessment of ways to optimally organize packets into streams for efficient parsing. Thorough filtering of events and event fields such that the target activity is concisely isolated, should be applied early in the processing stage to reduce resource utilization.

With respect to TCP/IP gateway traffic monitoring, we have investigated a variety of ways to categorize and isolate groups of packets from an arbitrary packet stream. Individual packet streams can be filtered based on different isolation criteria, such as

- *Discarded traffic:* packets not allowed through the gateway because they violate filtering rules [iii]
- *Pass-through traffic:* packets allowed into the internal network from external sources.
- *Protocol-specific traffic:* packets pertaining to a common protocol as designated in the packet header. One example is the stream of all ICMP packets that reach the gateway.
- *Unassigned port traffic:* packets targeting ports to which the administrator has not assigned any network service and that also remain unblocked by the firewall.
- *Transport management messages:* packets involving transport-layer connection establishment, control, and termination (e.g., TCP SYN, RESET, ACK, (window resize)).
- *Source-address monitoring:* packets whose source addresses match well-known external sites (e.g., connections from satellite offices) or have raised suspicion from other monitoring efforts.
- *Destination-address monitoring:* all packets whose destination addresses match a given internal host or workstation.
- *Application-layer monitoring:* packets targeting a particular network service or application. This stream isolation may translate to parsing packet headers for IP/port matches (assuming an established binding between port and service) and rebuilding datagrams.

In the following sections we discuss how such traffic streams can be statistically and heuristically analyzed to provide insight into malicious and erroneous external traffic. Alternative sources of event data are also available from the report logs produced by the various gateways, firewalls, routers, and proxy-servers (e.g., router syslogs can in fact be used to collect packet information from several products). We explore how statistical and signature analysis techniques can be employed to monitor various elements within TCP/IP event streams that flow through network gateways. We present specific techniques for detecting external entities that attempt to subvert or bypass internal network services. Techniques are suggested for detecting attacks against the underlying network infrastructure, including attacks using corruption or forgery of legitimate traffic in an attempt to negatively affect routing services, application-layer services, or other network controls. We suggest how to extend our surveillance techniques to recognize network faults and other exceptional activity. We also discuss issues of distributed result correlation.

4. Traffic Analysis with Statistical Anomaly Detection

SRI has been involved in statistical anomaly-detection research for over a decade [1], [5], [10]. Our previous work focused on the profiling of user activity through audit-trail analysis. Within the EMERALD project, we are extending the underlying statistical algorithms to profile various aspects of network traffic in search of response- or alert-worthy anomalies.

The statistical subsystem tracks subject activity via one or more variables called *measures*. The statistical algorithms employ four classes of measures: categorical, continuous, intensity, and event distribution. *Categorical* measures are those that assume values from a categorical set, such as originating host identity, destination host, and port number. *Continuous* measures are those for which observed values are numeric or ordinal, such as number of bytes transferred. Derived measures also track the intensity of activity (that is, the rate of events per unit time) and the "meta-distribution" of the measures affected by recent events. These derived measure types are referred to as *intensity* and *event distribution*.

The system we have developed maintains and updates a description of a subject's behavior with respect to these measure types in a compact, efficiently updated *profile*. The profile is subdivided into short- and long-term elements. The short-term profile accumulates values between updates, and exponentially ages values for comparison to the long-term profile. As a consequence of the aging mechanism, the short-term profile characterizes the recent activity of the subject, where "recent" is determined by the dynamically configurable aging parameters used. At update time (typically, a time of low system activity), the update function folds the short-term values observed since the last update into the long-term profile, and the short-term profile is cleared. The long-term profile is itself slowly aged to adapt to changes in subject activity. Anomaly scoring compares related attributes in the short-term profile against the long-term profile. As all evaluations are done against empirical distributions, no assumptions of parametric distributions are made, and multi-modal and categorical distributions are accommodated. Furthermore, the algorithms we have developed require no *a priori* knowledge of intrusive or exceptional activity. A more detailed mathematical description of these algorithms is given in [9], [26].

Our earlier work considered the subject class of users of a computer system and the corresponding event stream the system audit trail generated by user activity. Within the EMERALD project, we generalize these concepts so that components and software such as network gateways, proxies, and network services can themselves be made subject classes. The generated event streams are obtained from log files, packet analysis, and—where required—special-purpose instrumentation made for services of interest (e.g., FTP, HTTP, or SMTP). As appropriate, an event stream may be analyzed as a single subject, or as multiple subjects, and the same network activity can be analyzed in several ways. For example, an event stream of dropped packets permits analyses that track the reason each packet was rejected. Under such a scenario, the firewall rejecting the packet is the subject, and the measures of interest are the reason the packet was dropped (a categorical measure), and the rate of dropped packets in the recent past (one or more intensity measures tuned to time intervals of seconds to minutes). Alternatively, these dropped packets may be parsed in finer detail, supporting other analyses where the subject is, for

example, the identity of the originating host.

EMERALD can also choose to separately define satellite offices and "rest of world" as different subjects for the same event stream. That is, we expect distinctions from the satellite office's use of services and access to assets to deviate widely from sessions originating from external nonaffiliated sites. Through satellite session profiling, EMERALD can monitor traffic for signs of unusual activity. In the case of the FTP service, for example, each user who gives a login name is a subject, and "anonymous" is a subject as well. Another example of a subject is the network gateway itself, in which case there is only one subject. All subjects for the same event stream (that is, all subjects within a subject class) have the same measures defined in their profiles, but the internal profile values are different.

As we migrate our statistical algorithms that had previously focused on user audit trails with users as subjects, we generalize our ability to build more abstract profiles for varied types of activity captured within our generalized notion of an event stream. In the context of statistically analyzing TCP/IP traffic streams, profiling can be derived from a variety of traffic perspectives, including profiles of

- Protocol-specific transactions (e.g., all ICMP exchanges)
- Sessions between specific internal hosts and/or specific external sites
- Application-layer-specific sessions (e.g., anonymous FTP sessions profiled individually and/or collectively)
- Discarded traffic, measuring attributes such as volume and disposition of rejections
- Connection requests, errors, and unfiltered transmission rates and disposition

Event records are generated either as a result of activity or at periodic intervals. In our case, activity records are based on the content of IP packets or transport-layer datagrams. Our event filters also construct interval summary records, which contain accumulated network traffic statistics (at a minimum, number of packets and number of kilobytes transferred). These records are constructed at the end of each interval (e.g., once per N seconds).

EMERALD's statistical algorithm adjusts its short-term profile for the measure values observed on the event record. The distribution of recently observed values is evaluated against the long-term profile, and a distance between the two is obtained. The difference is compared to a historically adaptive, subject-specific deviation. The empirical distribution of this deviation is transformed to obtain a score for the event. Anomalous events are those whose scores exceed a historically adaptive, subject-specific score threshold based on the empirical score distribution. This nonparametric approach handles all measure types and makes no assumptions on the modality of the distribution for continuous measures.

The following sections provide example scenarios of exceptional network activity that can be measured by an EMERALD statistical engine deployed to network gateways.

4.1 Categorical Measures in Network Traffic

Categorical measures assume values from a discrete, nonordered set of possibilities. Examples of categorical measures include

- Source/destination address: One expects, for example, accesses from satellite offices to originate from a set of known host identities.
- Command issued: While any single command may not in itself be anomalous, some intrusion scenarios (such as "door-knob rattling") give rise to an unusual mix of commands in the short-term profile.
- Protocol: As with commands, a single request of a given protocol may not be anomalous, but an unusual mix of protocol requests, reflected in the short-term profile, may indicate an intrusion.
- Errors and privilege violations: We track the return code from a command as a categorical measure; we expect the distribution to reflect only a small percent of abnormal returns (the actual rate is learned in the long-term profile). While some rate of errors is normal, a high number of exceptions in the recent past is abnormal. This is reflected both in unusual frequencies for abnormal categories, detected here, and unusual count of abnormal returns, tracked as a continuous measure as described in Section 4.2.
- Malformed service requests: Categorical measures can track the occurrence of various forms of bad requests or malformed packets directed to a specific network service.
- Malformed packet disposition: Packets are dropped by a packet filter for a variety of reasons, many of which are innocuous (for example, badly formed packet header). Unusual patterns of packet rejection or error messages could lead to insight into problems in neighboring systems or more serious attempts by external sites to probe internal assets.
- File handles: Certain subjects (for example, anonymous FTP users) are restricted as to which files they can access. Attempts to access other files or to write read-only files appear anomalous. Such events are often detectable by signature analysis as well.

The statistical component builds empirical distributions of the category values encountered, even if the list of possible values is open-ended, and has mechanisms for "aging out" categories whose long-term probabilities drop below a threshold.

The following is an example of categorical measures used in the surveillance of proxies for services such as SMTP or FTP. Consider a typical data-exchange sequence between an external client and an internal server within the protected network. Anonymous FTP is restricted to certain files and directories; the names of these are categories for measures pertaining to file/directory reads and (if permitted) writes. Attempted accesses to unusual directories appear anomalous. Monitors dedicated to ports include a categorical measure whose values are the protocol used. Invalid requests often lead to an access violation error; the type of error associated with a request is another example of a categorical measure, and the count or rate of errors in the recent past is tracked as continuous measures, as described in Section

4.2 Continuous Measures in Network Traffic

Continuous measures assume values from a continuous or ordinal set. Examples include inter-event time (difference in time stamps between consecutive events from the same stream), counting measures such as the number of errors of a particular type observed in the recent past, and network traffic measures (number of packets and number of kilobytes). The statistical subsystem treats continuous measures by first allocating bins appropriate to the range of values of the underlying measure, and then tracking the frequency of observation of each value range. In this way, multi-modal distributions are accommodated and much of the computational machinery used for categorical measures is shared.

Continuous measures are useful not only for intrusion detection, but also support the monitoring of health and status of the network from the perspective of connectivity and throughput. An instantaneous measure of traffic volume maintained by a gateway monitor can detect a sudden and unexpected loss in the data rate of received packets, when this volume falls outside historical norms for the gateway. This sudden drop is specific both to the gateway (the subject, in this case) and to the time of day (e.g., the average sustained traffic rate for a major network artery is much different at 11:00 a.m. than at midnight).

In our example discussion of an FTP service in Section 4.1, attempts to access unallowed directories or files result in errors. The recently observed rate of such errors is continuously compared with the rate observed over similar time spans for other FTP sessions. Some low rate of error due to misspellings or innocent attempts is to be expected, and this would be reflected in the historical profile for these measures. An excess beyond historical norms indicates anomalous activity.

Continuous measures can also work in conjunction with categorical measures to detect excessive data transfers or file uploads, or excessive mail relaying, as well as excessive service-layer errors by external clients. Categorical and continuous measures have proven to be the most useful for anomaly detection in a variety of contexts.

We next describe the two derived measure types, *intensity* and *event distribution*, which detect anomalies related to recent traffic volume and the mix of measures affected by this traffic.

4.3 Measuring Network Traffic Intensity

Intensity measures distinguish whether a given volume of traffic appears consistent with historical observations. These measures reflect the intensity of the event stream (number of events per unit time) over time intervals that are tunable. Typically, we have defined three intensity measures per profile, which, with respect to user activity monitoring, were scaled at intervals of 60 seconds, 600 seconds, and 1 hour. Applied to raw event streams, intensity measures are particularly suited for detecting flooding attacks, while also providing insight into other anomalies.

EMERALD uses volume analyses to help detect the introduction of malicious traffic, such as traffic intended to cause service denials or perform intelligence gathering, where such traffic may not necessarily be violating filtering policies. A sharp increase in the overall volume of discarded packets, as well as analysis of the disposition of the discarded packets (as discussed in Section 4.1), can provide insight into unintentionally malformed packets resulting from poor line quality or internal errors in neighboring hosts. High volumes of discarded packets can also indicate more maliciously intended transmissions such as scanning of UDP ports or IP address scanning via ICMP echoes. Excessive numbers of mail expansion requests (XMRs) may indicate intelligence gathering, perhaps by spammers. These and other application-layer forms of doorknob rattling can be detected by an EMERALD statistical engine when filtering is not desired.

Alternatively, a sharp increase in events viewed across longer durations may provide insight into a consistent effort to limit or prevent successful traffic flow. Intensity measures of transport-layer connection requests, such as a volume analysis of SYN-RST messages, could indicate the occurrence of a SYN-attack [17] against port availability (or possibly for port scanning). Variants of this could include intensity measures of TCP/FIN messages [14], considered a more stealthy form of port scanning.

Monitoring overall traffic volume and bursty events by using both intensity and continuous measures provides some interesting advantages over other monitoring approaches, such as user-definable heuristic rules that specify fixed thresholds. In particular, the intensity of events over a duration is relative in the sense that the term "high volume" may reasonably be considered different at midnight than at 11:00 a.m. The notion of high bursts of events might similarly be unique to the role of the target system in the intranet (e.g., web server host versus a user workstation). Rule developers would need to carefully define thresholds based on many factors unique to the target system. On the other hand, the statistical algorithms would, over time, build a target-specific profile that could evaluate event intensity for the given system over a variety of time slices such as the time of day (e.g., business hours versus afterhours) and/or day of the week (e.g., weekday versus weekend).

4.4 Event Distribution Measures

The event-distribution measure is a meta-measure that monitors which other measures in the profile are affected by each event. For example, an *ls* command in an FTP session affects the directory measure, but does not affect measures related to file transfer. This measure is not interesting for all event streams. For example, all network-traffic event records affect the same measures (number of packets and kilobytes) defined for that event stream, so the event distribution does not change.

On the other hand, event-distribution measures are useful in correlative analysis achieved via the "Monitor of Monitors" approach. Here, each monitor contributes to an aggregate event stream for the domain of the correlation monitor. These events are generated only when the individual monitor decides that the recent behavior is anomalous (though perhaps not sufficiently anomalous by itself to trigger a declaration). Measures recorded include time stamp, monitor identifier, subject

identifier, and measure identities of the most outlying measures. Overall intensity of this event stream may be indicative of a correlated attack. The distribution of which monitors and which measures are anomalous is likely to be different with an intrusion or malfunction than with the normal "innocent exception." (See Section 6 for a further discussion on result correlation.)

4.5 Statistical Session Analysis

Statistical anomaly detection via the methods described above enables EMERALD to answer questions such as how the current anonymous FTP session compares to the historical profile of all previous anonymous FTP sessions. Mail exchange could be similarly monitored for atypical exchanges (e.g., excessive mail relays).

Continuing with the example of FTP, we assign FTP-related events to a subject (the login user or "anonymous"). As several sessions may be interleaved, we maintain separate short-term profiles for each, but may score against a common long-term profile (for example, short-term profiles are maintained for each "anonymous" FTP session, but each is scored against the historical profile of "anonymous" FTP sessions). The aging mechanism in the statistics module allows it to monitor events either as the events occur or at the end of the session. We have chosen the former approach (analyze events as they happen), as it potentially detects anomalous activity in a session before that session is concluded.

5. Traffic Analyzing with Signature Analysis

Signature analysis is a process whereby an event stream is mapped against abstract representations of event sequences known to indicate the target activity of interest. Signature engines are essentially expert systems whose rules fire as event records are parsed that appear to indicate suspicious, if not illegal, activity. Signature rules may recognize single events that by themselves represent significant danger to the system, or they may be chained together to recognize sequences of events that represent an entire penetration scenario.

However, simplistic event-to-rule binding alone does not necessarily provide enough indication to ensure accurate detection of the target activity. Signature analyses must also distinguish whether an event sequence being witnessed is actually transitioning the system into the anticipated compromised state. In addition, determining whether a given event sequence is indicative of an attack may be a function of the preconditions under which the event sequence is performed. Example coding schemes for representing operating system penetrations through audit trail analysis are [12], [18], [19].

Using basic signature-analysis concepts, EMERALD can support a variety of analyses involving packet and transport datagrams as event streams. For example, address spoofing, tunneling, source routing [21], SATAN [27] attack detection,

and abuse of ICMP messages (Redirect and Destination Unreachable) messages in particular) [4] could all be encoded and detected by signature engines that guard network gateways. The heuristics for analyzing headers and application datagrams for some of these abuses are not far from what is already captured by some filtering tools. In fact, it is somewhat difficult to justify the expense of passively monitoring the traffic stream for such activity when one could turn such knowledge into filtering rules.[iv]

Regardless, there still remain several examples that help justify the expense of employing signature analyses to monitor network traffic. In particular, there are points where the appearance of certain types of legitimate traffic introduces questions regarding the motives of the traffic source. Distinguishing benign requests from illicit ones may be fairly difficult, and such questions are ultimately site-specific. For example, EMERALD surveillance modules can encode thresholds to monitor activity such as the number of fingers, pings, or failed login requests to accounts such as guest, demo, visitor, anonymous FTP, or employees who have departed the company. Threshold analysis is a rudimentary, inexpensive technique that records the occurrence of specific events and, as the name implies, detects when the number of occurrences of that event surpasses a reasonable count.

In addition, we are developing heuristics to support the processing of application-layer transactions derived from packet monitoring. EMERALD's signature analysis module can sweep the data portion of packets in search of a variety of transactions that indicate suspicious, if not malicious, intentions by the external client. While traffic filtering rules may allow external traffic through to an internally available network service, signature analysis offers an ability to model and detect transaction requests or request parameters, alone or in combination, that are indicative of attempts to maliciously subvert or abuse the internal service. EMERALD's signature engine, for example, is capable of real-time parsing of FTP traffic through the firewall or router for unwanted transfers of configuration or specific system data, or anonymous requests to access non-public portions of the directory structure. Similarly, EMERALD can analyze anonymous FTP sessions to ensure that the file retrievals and uploads/modifications are limited to specific directories. Additionally, EMERALD's signature analysis capability is being extended to session analyses of complex and dangerous, but highly useful, services like HTTP or Gopher.

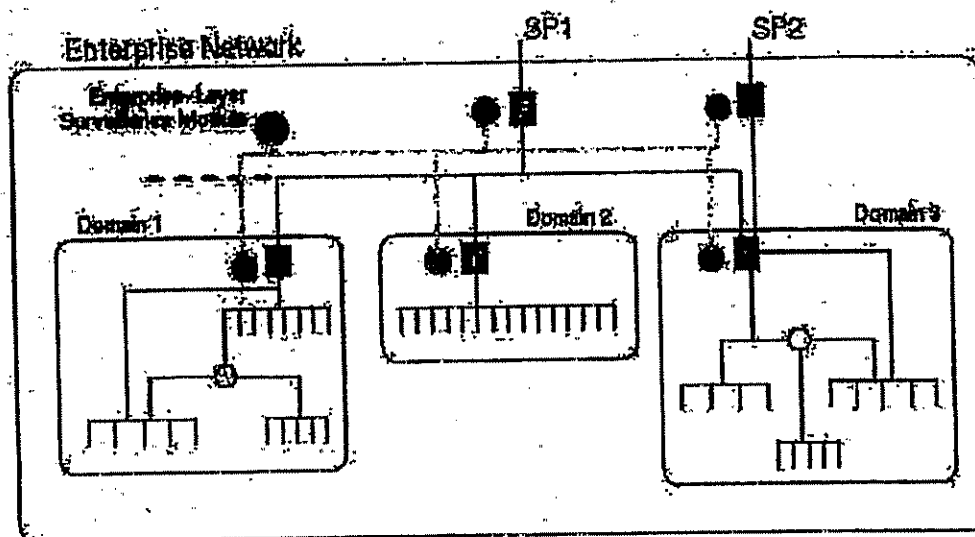
Another interesting application of signature analysis is the scanning of traffic directed at high-numbered unused ports (i.e., ports to which the administrator has not assigned a network service). Here, datagram parsing can be used to study network traffic after some threshold volume of traffic, directed at an unused port, has been exceeded. A signature module can employ a knowledge base of known telltale datagrams that are indicative of well-known network-service protocol traffic (e.g., FTP, Telnet, SMTP, HTTP). The signature module then determines whether the unknown port traffic matches any known datagram acts. Such comparisons could lead to the discovery of network services that have been installed without an administrator's knowledge.

6. Composable Surveillance of Network Traffic

The focus of surveillance need not be limited to the analysis of traffic streams through a single gateway. An extremely useful extension of anomaly detection and signature analyses is to support the hierarchical correlation of analysis results produced by multiple distributed gateway surveillance modules. Within the EMERALD framework, we are developing meta-surveillance modules that analyze the anomaly and signature reports produced by individual traffic monitors dispersed to the various entry points of external traffic into local network domains.

This concept is illustrated in Figure 1, which depicts an example enterprise network consisting of interconnected local network domains.^[v] These local domains are independently administered, and could perhaps correspond to the division of computing assets among departments within commercial organizations or independent laboratories within research organizations. In this figure, connectivity with the external world is provided through one or more service providers (SP1 and SP2), which may provide a limited degree of filtering based on source address (to avoid address spoofing), as well as other primitive checks such as monitoring checksum.

Example Network Deployment of Surveillance Monitors



Inside the perimeter of the enterprise, each local domain maintains its traffic filtering control (F-boxes) over its own subnetworks. These filters enforce domain-specific restriction over issues such as UDP port availability, as well as acceptable protocol traffic. EMERALD surveillance monitors are represented by the S-circles, and are deployed to the various entry points of the enterprise and domains.

EMERALD surveillance modules develop analysis results that are then directed up to an enterprise-layer monitor, which correlates the distributed results into a meta-event stream. The enterprise monitor is identical to the individual gateway

monitors (i.e., they use the same code base), except that it is configured to correlate activity reports produced by the gateway monitors. The enterprise monitor employs both statistical anomaly detection and signature analyses to further analyze the results produced by the distributed gateway surveillance modules, searching for commonalities or trends in the distributed analysis results.

The following sections focus on aggregate analyses that may induce both local response and/or enterprise-wide response. We enumerate some of the possible ways that analysis results from the various surveillance modules can be correlated to provide insight into more global problems not visible from the narrow perspective of local entry-point monitoring.

6.1 Commonalities among Results

One issue of direct interest is whether there exist commonalities in analysis results across surveillance modules that are examining mutually exclusive event streams. For example, a scenario previously discussed was that of a statistical engine observing a drastic increase in the number of discarded packets at the entry point to a domain, perhaps even observing the majority cause for packet discards. Depending on the degree of increase, a local domain administrator could be persuaded to take actions to help alleviate or remove the cause of the failed packets. However, if on a given day all such domains throughout the enterprise similarly observed marked increases in discarded packet volume, the response could propagate from being a local concern to being an enterprise-wide issue. Similarly, commonalities across domains in excessive levels of protocol-specific errors or signature engines detecting unwanted activity across multiple domains could lead to enterprise-layer responses.

We might also choose to distinguish excessive types of certain traffic in an effort to check for intelligence gathering by outsiders who submit requests such as finger, echo, or mail alias expansion, to multiple domains in the enterprise (i.e., round-robin doorknob rattling). The objective of such a technique might be to avoid detection from both local network intensity and/or continuous measures by spreading out the probes to multiple independently monitored domains. Through aggregate analysis, we could maintain the enterprise-wide profile of probes of this type, and detect when an unusual number or mix of these probes occurs. While such probes may not appear excessive from the local domain perspective, the enterprise overall may observe a marked increase worthy of response.

In addition, we can add a layer of traffic-rate monitoring by profiling the overall volume of enterprise traffic expected throughout various slices of the day and week. Local monitors may use continuous measures to detect drastic declines in packet volumes that could indicate transmission loss or serious degradation. However, it is conceivable that the degradation from the local domain perspective, while significant, is not drastic enough to warrant active response. At the same time, we may find through results correlation that the aggregate of all domains producing reports of transmission rate degradation during the same time period could warrant attention at the enterprise layer. Thus, local domain activity below the severity of

warranting a response could in aggregation with other activity be found to warrant a response.

6.2 Sequential Trend Analysis

Of general use to meta-surveillance is the modeling of activity for sequential trends in the appearance of problematic traffic. For example, this could entail correlating the analyses of local monitors, looking for trends in the propagation of application-layer datagrams for error or ICMP packets. While local responses to error messages could be handled by the local domain administrators, reports of errors spreading across all domains might more effectively be addressed by those responsible for connections between the enterprise and the service provider.

Attacks repeated against the same network service across multiple domains can also be detected through enterprise-layer correlation. For example, multiple surveillance modules deployed to various local domains in the enterprise might begin to report, in series, suspicious activity observed within sessions employing the same network service. Such reports could lead to enterprise-layer responses or warnings to other domains that have not yet experienced or reported the session anomalies. In this sense, results correlation enables the detection of spreading attacks against a common service, which first raise alarms in one domain, and gradually spread domain by domain to affect operations across the enterprise.

We are studying the use of fault-relationship models [22], in which recognition of a problem in one network component (e.g., loss of connectivity or responsiveness) could propagate as different problems in neighboring hosts (e.g., buffer overflows or connection timeout due to overloads). Our enterprise monitor employs rule-based heuristics to capture such relationship models.

7. Response Handling

Once a problem is detected, the next challenge is to formulate an effective response. In many situations, the most effective response may be no response at all, in that every response imposes some cost in system performance or (worse) human time. The extent to which a decision unit contains logic to filter out uninteresting analysis results may mean the difference between effective monitoring units and unmanageable (soon to be disabled) monitoring units. For certain analysis results such as the detection of known hostile activity through signature analyses, the necessity for response invocation may be obvious. For other analysis results such as anomaly reports, response units may require greater sophistication in the invocation logic.

Fundamental to effective response handling is the accurate identification of the source responsible for the problem. However, unlike audit-trail analysis where event-record fields such as the subject ID are produced by the OS kernel, attackers have

direct control over the content and format of packet streams. Packet forgery is straightforward, and one must take care to avoid allowing attackers to manipulate response logic to harm legitimate user connectivity or cause service denials throughout the network. Some techniques have been proposed to help track network activity to the source [24].

Another issue is how to tailor a response that is appropriate given the severity of the problem, and that provides a singular effect to address the problem without harming the flow of legitimate network traffic. Countermeasures range from very passive responses, such as passive results dissemination, to highly aggressive actions, such as severing a communication channel. Within EMERALD, our response capabilities will employ the following general forms of response:

- **Passive results dissemination:** EMERALD monitors can make their analysis results available for administrative review. We are currently exploring techniques to facilitate passive dissemination of analysis results by using already-existing network protocols such as SNMP, including the translation of analysis results into an intrusion-detection management information base (MIB) structure. However, whereas it is extremely useful to integrate results dissemination into an already-existing infrastructure, we must balance this utility with the need to preserve the security and integrity of analysis results.
- **Assertive results dissemination:** Analysis results can be actively disseminated as administrative alerts. While the automatic dissemination of alerts may help to provide timely review of problems by administrators, this approach may be the most expensive form of response, in that it requires human oversight [vi].
- **Dynamic controls over logging configuration:** EMERALD monitors can perform limited control over the (re)configuration of logging facilities within network components (e.g., routers, firewalls, network services, audit daemons).
- **Integrity checking probes:** EMERALD monitors may invoke handlers that validate the integrity of network services or other assets. Integrity probes may be particularly useful for ensuring that privileged network services have not been subverted [vii].
- **Reverse probing:** EMERALD monitors may invoke probes in an attempt to gather as much counterintelligence about the source of suspicious traffic by using features such as *traceroute* or *finger*. However, care is required in performing such actions, as discussed in [4].
- **Active channel termination:** An EMERALD monitor can actively terminate a channel session if it detects specific known hostile activity. This is perhaps the most severe response, and care must be taken to ensure that attackers do not manipulate the surveillance monitor to deny legitimate access.

8. Conclusion

We have described event-analysis techniques developed in the intrusion-detection community, and discussed their application to monitoring TCP/IP packet streams through network gateways. We present a variety of exceptional activity (both malicious and nonmalicious) to which these analysis techniques could be applied. Table 1 summarizes the analyzable exceptional network activity presented in this paper, and identifies which method (statistical anomaly detection, signature analysis, or hierarchical correlation) can be utilized to detect the activity.

These examples help to justify the expense of gateway surveillance monitors, even in the presence of sophisticated traffic-filtering mechanisms. Indeed, several of the example forms of "interesting traffic" listed in Table 1 are not easily, if at all, preventable using filtering mechanisms. In addition, our surveillance modules may even help to tune or point out mistakes in filtering rules that could lead to the accidental discarding of legitimate traffic. The surveillance modules may detect the occurrence of traffic that appears to be anomalous or abusive, regardless of whether the traffic is allowed to enter, or is prevented from entering, through the network gateway. Furthermore, these techniques may extend to nonmalicious problem detection such as failures in neighboring systems.

While this paper is intended to justify and illustrate the complementary nature of combining surveillance capabilities with filtering mechanisms, in future research we will explore the practical aspects of monitor deployment, including performance analysis and secure integration into supporting network infrastructure (e.g., network management). Perhaps even more than traditional audit-based intrusion-detection developers, network monitor developers must carefully assess the optimum ways to organize and isolate the relevant traffic from which their analyses are based. The added dimension of control and insight into network operations gained by well-integrated surveillance modules is well worth consideration.

Analysis Description	Stat. Categ. Meas.	Stat. Conti. Meas.	Stat. Inten. Meas.	Sign. Analy.	Hier. Corr.
Protocol-specific anomalies such as excessive data transfers (FTP uploads, email relays, other large data transfers)	X	X	X		
Port/service misses, including excessive errors or unknown command exchanges	X		X		
Discarded packet volume			X		
Discarded packet disposition (analysis of rejection patterns)	X	X			
Excessive transport-layer connection requests, including heavy syn-ack message usage	X		X		
Anonymous session comparisons against historical usage	X	X	X		

Satellite office profiling	X	X	X		
Sudden drops or floods in data rate (specific to system, time of day, day of week, and so forth)		X	X		
Address/port scanning and other general doorknob rattling			X		
Excessive drops in line quality compared to historical quality		X	X		
Detection of filterable events (e.g., ICMP message abuse, address spoofing, tunnelling, source/port routing, BATAN signatures)				X	
Event thresholds for events reflecting site-specific concerns				X	
Detection of user-installed network services on unregistered ports				X	
Packet data sweeps for application-layer probes, looking for troublesome data transfers or requests				X	
Aggregate analysis across the enterprise for round-robin doorknob rattling that attempts to defeat domain-layer intensity measures					X
Aggregate analysis of low-level degradation of services or throughput across the enterprise					X
Trend analysis for error propagation occurring across multiple domains					X
Spreading attacks that may indicate worm or fault interrelationships among network modules					X

Endnotes

- i. Recent product examples, such as ASIM and Net Ranger, that follow the passive packet monitoring approach have since gained wide deployment in some Department of Defense network facilities.
- ii. We use the terms *enterprise* and *intranet* interchangeably; both exist ultimately as cooperative communities of independently administered domains, communicating together with supportive network infrastructure such as firewalls, routers, and bridges.
- iii. Of particular added value in assessing this traffic would be some indication of why a given packet was rejected. A generic solution for deriving this *disposition* information without dependencies on the firewall or router is difficult. Such information would be a useful enhancement to packet-rejection handlers.
- iv. On the other hand, one may also suggest a certain utility in simply having real-time mechanisms to detect, report, and hierarchically correlate attempts by external sources to forward undesirable packets through a gateway.
- v. This is one example network filtering strategy that is useful for illustrating result correlation. Other strategies are possible.

- vi. Consider a network environment that on average supports 100,000 external transactions (the definition of transaction is analysis-target-specific) per day. Even if only 0.1% of the transactions were found worthy of administrative review, administrators would be asked to review 100 transactions a day.
- vii. A significant number of network attacks target the subversion of privileged network service. CERT Advisories CA-97.16, CA-97.12, CA-97.05 give a few recent examples.

References

1. D.Anderson, T.Frivoold, and A.Valdes. Next-generation intrusion-detection expert system (NIDES): Final technical report. *Technical report*, Computer Science Laboratory, SRI International, Menlo Park, CA, 16 November 1994.
2. B.Chapman and B.Zwicky. *Building Internet firewalls*. O'Reilly and Associates, Inc. Sebastopol, CA, 1995.
3. D.Chapman. Network (in)security through IP packet filtering. In *Proceedings of the Third USENIX Unix Security Symposium*, Baltimore, MD, September 1992.
4. W.R. Cheswick and S.M. Bellovin. *Firewalls and internet security: Repelling the wily hacker*. Addison-Wesley, Reading, MA, 1994.
5. D.E. Denning. An intrusion-detection model. *IEEE Transactions on Software Engineering*, 13(2), February 1987.
6. L.T. Heberlein, G.Dias, K.N. Levitt, B.Mukherjee, J.Wood, and D.Wolber. A network security monitor. In *Proceedings of the 1990 Symposium on Research in Security and Privacy*, pages 296-303, Oakland, CA, May 1990. IEEE Computer Society.
7. K.Jackson, D.DuBois, and C.Stallings. An expert system application for network intrusion detection. In *Proceedings of the Fourteenth Computer Security Group Conference*. Department of Energy, 1991.
8. G.Jakobson and M.D. Weissman. Alarm correlation. *IEEE Network*, pages 52-59, November 1993.
9. H.S. Javitz and A.Valdes. The NIDES statistical component description and justification. Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA, March 1994.
10. H.S. Javitz, A.Valdes, D.E. Denning, and P.G. Neumann. Analytical techniques development for a statistical intrusion-

- detection system (SIDS) based on accounting records. *Technical report*, SRI International, Menlo Park, CA, July 1986.
11. S.Kliger, S.Yemini, Y.Yemini, D.Ohsia, and S.Stolfo. A coding approach to event correlation. In *Proceedings of the Fourth International Symposium on Integrated Network Management (IFIP/IEEE)*, Santa Barbara, CA, pages 266-277. Chapman and Hall, London, England, May 1995.
 12. T.F. Lunt, R.Jagannathan, R.Lee, A.Whitehurst, and S.Listgarten. Knowledge-based intrusion detection. In *Proceedings of the 1989 AI Systems in Government Conference*, March 1989.
 13. T.F. Lunt, A.Tamaru, F.Gilham, R.Jagannathan, C.Jalali, P.G. Neumann, H.S. Javitz, and A.Valdes. A real-time intrusion-detection expert system (IDES). *Technical report*, Computer Science Laboratory, SRI International, Menlo Park, CA, 28 February 1992.
 14. Uriel Maimon. Port scanning without the SYN flag. *Phrack Magazine*, vol. 7, issue 49.
 15. M.Mansouri-Samani and M.Sloman. Monitoring distributed systems. *IEEE Network*, pages 20-30, November 1993.
 16. K.Moyer, M.Bringer, J.Betser, C.Sunshine, G.Goldschmidt, and Y.Yemini. Decentralizing control and intelligence in network management. In *Proceedings of the Fourth International Symposium on Integrated Network Management (IFIP/IEEE)*, Santa Barbara, CA, pages 4-16. Chapman and Hall, London, England, May 1995.
 17. Robert T. Morris. A weakness in the 4.2BSD UNIX TCP/IP software. In *Computing Science Technical Report 117*. AT&T Bell Laboratories, Murray Hills, NJ, 25 February 1985.
 18. A.Mounji, B.Le Charlier, and D.Zampunieria. Distributed audit trail analysis. In *Proceedings of the ISOC 1995 Symposium on Network and Distributed System Security*, pages 102-112, February 1995.
 19. P.A. Porras. STAT: A State Transition Analysis Tool for intrusion detection. Master's thesis, Computer Science Department, University of California, Santa Barbara, July 1992.
 20. P.A. Porras and P.G. Neumann. EMERALD: Event monitoring enabling responses to anomalous live disturbances. In *National Information Systems Security Conference*, pages 353-365, Baltimore, MD, October 1997.
 21. J.Postel. Internet protocol, request for comment, RFC 791. *Technical report*, Information Sciences Institute, September 1981.

22. L. Ricciulli and N. Shacham. Modeling correlated alarms in network management systems. In *Communication Networks and Distributed Systems Modeling and Simulation*, 1997.
23. S.R. Snapp, J. Brentano, G.V. Dias, T.L. Goan, L.T. Heberlein, C.-L. Ho, K.N. Levitt, B. Mukherjee, S. Smaha, T. Gramsc, D.M. Teal, and D. Mansur. DIDS (Distributed Intrusion Detection System)—motivation, architecture, and an early prototype. In *Proceedings of the Fourteenth National Computer Security Conference*, pages 167–176, Washington, D.C., 1–4 October 1991. NIST/NCSC.
24. S. Staniford-Chen, S. Cheung, R. Crawford, M. Dilger, J. Frank, J. Hogland, K. Levitt, C. Woo, R. Yip, and D. Zeckle. GridS—a graph based intrusion detection system for large networks. In *Proceedings of the Nineteenth National Information Systems Security Conference*, pages 361–370 (Volume I), Washington, D.C., October 1996. NIST/NCSC.
25. S. Staniford-Chen and L.T. Heberlein. Holding intruders accountable on the internet. In *Proceedings of the IEEE Symposium on Security and Privacy*, 1995.
26. A. Valdes and D. Anderson. Statistical methods for computer usage anomaly detection using NIDES. *Proceedings of the Third International Workshop on Rough Sets and Soft Computing (RSSC 94)*, San Jose, January 1995.
27. W. Venema. Project SATAN: UNIX/internet security. In *Proceedings of the COMSEC-95 Conference*, Elsevier, London, 1995.

EXHIBIT C

EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances †

Phillip A. Porras and Peter G. Neumann
porras@csl.sri.com and neumann@csl.sri.com
<http://www.csl.sri.com/intrusion.html>

Computer Science Laboratory
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025-3493

Abstract— The EMERALD (Event Monitoring Enabling Responses to Anomalous Live Disturbances) environment is a distributed scalable tool suite for tracking malicious activity through and across large networks. EMERALD introduces a highly distributed, building-block approach to network surveillance, attack isolation, and automated response. It combines models from research in distributed high-volume event-correlation methodologies with over a decade of intrusion detection research and engineering experience. The approach is novel in its use of highly distributed, independently tunable, surveillance and response monitors that are deployable polymorphically at various abstract layers in a large network. These monitors contribute to a streamlined event-analysis system that combines signature analysis with statistical profiling to provide localized real-time protection of the most widely used network services on the Internet. Equally important, EMERALD introduces a recursive framework for coordinating the dissemination of analyses from the distributed monitors to provide a global detection and response capability that can counter attacks occurring across an entire network enterprise. Further, EMERALD introduces a versatile application programmers' interface that enhances its ability to integrate with heterogeneous target hosts and provides a high degree of interoperability with third-party tool suites.

Keywords— Network security, intrusion detection, coordinated attacks, anomaly detection, misuse detection, information warfare, system survivability, insider threat, outsider threat.

I. INTRODUCTION

Our infrastructures of highly integrated information systems, both military and commercial, have become one of the key assets on which we depend for competitive advantage. These information infrastructures tend to be conglomerates of integrated commercial-off-

the-shelf (COTS) and non-COTS components, interoperating and sharing information at increasing levels of demand and capacity. These systems are relied on to manage a growing list of needs including transportation, commerce, energy management, communications, and defense.

Unfortunately, the very interoperability and sophisticated integration of technology that make our infrastructures such valuable assets also make them vulnerable to attack, and make our dependence on our infrastructures a potential liability. We have had ample opportunity to consider numerous examples of vulnerabilities and attacks against our infrastructures and the systems that use them. Attacks such as the Internet worm [21], [23] have shown us how our interconnectivity across large domains can be used against us to spread malicious code. Accidental outages such as the 1980 ARPAnet collapse [22] and the 1990 AT&T collapse [17] illustrate how seemingly localized triggering events can have globally disastrous effects on widely distributed systems. In addition, we have witnessed organized groups of miscreants [11], [17], local and foreign, performing malicious and coordinated attacks against varieties of online targets. We are keenly aware of the recurring examples of vulnerabilities that exist pervasively in network services, protocols, and operating systems, throughout our military and commercial network infrastructures. Even the deployment of newer more robust technologies does not fully compensate for the vulnerabilities in the multitude of legacy systems with which the newer systems must interoperate.

Yet, despite these examples, there remain no widely available robust tools to allow us to track malicious activity through and across large networks. The need for scalable network-aware surveillance and response technologies continues to grow.

† The work described here is currently funded by DARPA/ITO under contract number F30602-96-C-0294.

II. CHALLENGES TO SCALABLE NETWORK MISUSE DETECTION

As dependence on our network infrastructures continues to grow, so too grows our need to ensure the survivability of these assets. Investments into scalable network intrusion detection¹ will over time offer an important additional dimension to the survivability of our infrastructures. Mechanisms are needed to provide real-time detection of patterns in network operations that may indicate anomalous or malicious activity, and to respond to this activity through automated countermeasures. In addition, these mechanisms should also support the pursuit of individuals responsible for malicious activity through the collection and correlation of event data.

The typical target environment of the EMERALD project is a large enterprise network with thousands of users connected in a federation of independent administrative domains. Each administrative domain is viewed as a collection of local and network services that provide an interface for requests from individuals internal and external to the domain. Network services include features common to many network operating systems such as mail, HTTP, FTP, remote login, network file systems, finger, Kerberos, and SNMP. Some domains may share trust relationships with other domains (either peer-to-peer or hierarchical). Other domains may operate in complete mistrust of all others, providing outgoing connections only, perhaps severely restricting incoming connections. Users may be local to a single domain or may possess accounts on multiple domains that allow them to freely establish connections throughout the enterprise.

In the environment of an enterprise network, well-established concepts in computer security such as the *reference monitor* [3] do not apply well. A large enterprise network is a dynamic cooperative of interconnected heterogeneous systems that often exists more through co-dependence than hierarchical structure. Defining a single security policy over such an enterprise, let alone a single point of authority, is often not practical.

With traditional approaches to security being difficult to apply to network infrastructures in the large, the need to ensure survivability of these infrastructures raises important questions. One such question is, "*Can we build surveillance and response capabilities that can scale to very large enterprise networks?*" To do so will require us to overcome a number of challenges in cur-

rent intrusion-detection designs, many of which derive from the centralized paradigm of current architectures. While a fully distributed architecture could address some of these challenges, it too introduces tradeoffs in capabilities and performance. The following briefly summarizes challenges that exist in scaling intrusion-detection tools to large networks.

- **Event Generation and Storage:** Audit generation and storage has tended to be a centralized activity, and often gathers excessive amounts of information at inappropriate layers of abstraction. Centralized audit mechanisms place a heavy burden on the CPU and I/O throughput, and simply do not scale well with large user populations. In addition, it is difficult to extend centralized audit mechanisms to cover spatially distributed components such as network infrastructure (e.g., routers, filters, DNS, firewalls) or various common network services.

- **State-space Management and Rule Complexity:** In signature-based analyses, rule complexity can have a direct tradeoff with performance. A sophisticated rule structure able to represent complex/multiple event orderings with elaborate pre- or post-conditions may allow for very concise and well-structured penetration definitions. However, sophisticated rule structures may also impose heavy burdens in maintaining greater amounts of state information throughout the analysis, limiting their scalability to environments with high volumes of events. Shorter and simpler rules may impose lesser analysis and state-management burdens, helping to provide greater scalability and efficiency in event analysis. When speed is the key issue, the ultimate rule-set is one with no state-management needs — requiring no ordering and no time-consuming pre- and post-conditions to evaluate as events are processed. Simpler rules, however, also limit expressibility in misuse definitions, and can lead to inflated rule-bases to compensate for a single complex rule-set that might cover many variations of an attack. Clearly, there exists a tradeoff between highly complex and expressibly rich rule models versus shorter and simpler rules that individually require minimal state-management and analysis burdens.

- **Knowledge Repositories:** Expert systems separate their base of knowledge (rules of inference and state information regarding the target system) from both their analysis code and response logic in an effort to add to their overall modularity. There is some advantage to maintaining this knowledge base in a centrally located repository. Dynamic modification and control over this information is made easier when only single repositories need be modified. A centrally located knowledge repository is efficient for making pluggable rule-sets that add

¹ In this paper, the term "intrusion" is used broadly to encompass misuse, anomalies, service denials, and other deviations from acceptable system behavior.

to the generality and portability of the tool. However, in a highly distributed and high-volume event environment, a single repository combined with a single analysis engine can act as a choke-point. It also provides a single point of failure should the repository become unavailable or tainted.

• **Inference Architectures:** At the core of many signature-based expert systems exists an algorithm for accepting the input (in our case activity logs) and, based on a set of inference rules, directing the search for new information. This inference-engine model is very centralized in nature. In a large network, events and data flow asynchronously throughout the network in parallel and in volumes beyond what any centralized analysis technologies can process. A central analysis requires centralized collection of event information, and imposes the full burden (I/O, processing, and memory) of the analysis on those components on which the inference engine resides. This single-point-of-analysis model does not scale well. A completely distributed analysis, however, introduces its own challenges. Both global correlation and intelligent coordination among distributed analysis units impose significant resource overhead. Finding the optimal analysis paradigm between the continuum of the centralized expert-system approach and a fully decentralized analysis scheme is a key challenge in building a scalable inference architecture.

The physical and logical dispersion of the interfaces and controls among target systems and networks must be accommodated by the architecture of the distributed analysis system. Centralized intrusion-detection architectures deployed in highly distributed network environments experience difficulty in integrating and scaling their analysis paradigms to such environments. (Several of these issues are explored in [16]). The issues and limitation discussed above represent challenges to the very design and engineering assumptions on which much of the current intrusion-detection research is based.

The objective of the EMERALD work is to bring a collection of research and prototype development efforts into the practical world, in such a way that the analysis tools for detecting and interpreting anomalies and misuses can be applied and integrated into realistic network computing environments. The EMERALD project provides a critical step in demonstrating how to construct scalable and computationally realistic intrusion-detection mechanisms to track malicious activity within and across large networks. To do this, EMERALD employs detection and response components that are smaller and more distributed than previous intrusion-detection efforts, and that interoperate to provide composable surveillance.

EMERALD represents a significant departure from previous centralized host-based, user-oriented, intrusion-detection efforts that suffer poor scalability and integration into large networks. EMERALD's analysis scheme targets the external threat agent who attempts to subvert or bypass a domain's network interfaces and control mechanisms to gain unauthorized access to domain resources or prevent the availability of these resources. EMERALD employs a building-block architectural strategy using independent distributed surveillance monitors that can analyze and respond to malicious activity on local targets, and can interoperate to form an analysis hierarchy. This layered analysis hierarchy provides a framework for the recognition of more global threats to interdomain connectivity, including coordinated attempts to infiltrate or destroy connectivity across an entire network enterprise. Section III presents an architectural overview of EMERALD, and Section IV discusses its integration into distributed computing environments.

III. THE EMERALD NETWORK INTRUSION DETECTION ARCHITECTURE

EMERALD introduces a hierarchically layered approach to network surveillance that includes *service analysis* covering the misuse of individual components and network services within the boundary of a single domain; *domain-wide analysis* covering misuse visible across multiple services and components; and *enterprise-wide analysis* covering coordinated misuse across multiple domains. The objective of the service analysis is to streamline and decentralize the surveillance of a domain's network interfaces for activity that may indicate misuse or significant anomalies in operation. We introduce the concept of dynamically deployable, highly distributed, and independently tunable *service monitors*. Service monitors are dynamically deployed within a domain to provide localized real-time analysis of infrastructure (e.g., routers or gateways) and services (privileged subsystems with network interfaces). Service monitors may interact with their environment passively (reading activity logs) or actively via probing to supplement normal event gathering. This localized coverage of network services and domain infrastructure forms the lowest tier in EMERALD's layered network-monitoring scheme.

Information correlated by a service monitor can be disseminated to other EMERALD monitors through a *subscription-based* communication scheme. Subscription provides EMERALD's message system both a push and pull data exchange capability between monitor interoperation (see Section III-F). EMERALD client monitors are able to subscribe to receive the analysis

results that are produced by server monitors. As a monitor produces analysis results, it is then able to disseminate these results asynchronously to its client subscribers. Through subscription, EMERALD monitors distributed throughout a large network are able to efficiently disseminate reports of malicious activity without requiring the overhead of synchronous polling.

Domain-wide analysis forms the second tier of EMERALD's layered network surveillance scheme. A *domain monitor* is responsible for surveillance over all or part of the domain. *Domain monitors* correlate intrusion reports disseminated by individual service monitors, providing a domain-wide perspective of malicious activity (or patterns of activity). In addition to domain surveillance, the domain monitor is responsible for re-configuring system parameters, interfacing with other monitors beyond the domain, and reporting threats against the domain to administrators.

Lastly, EMERALD enables enterprise-wide analysis, providing a global abstraction of the cooperative community of domains. Enterprise-layer monitors correlate activity reports produced across the set of monitored domains. Enterprise-layer monitors focus on network-wide threats such as Internet worm-like attacks, attacks repeated against common network services across domains, and coordinated attacks from multiple domains against a single domain. Through this correlation and sharing of analysis results, reports of problems found by one monitor may propagate to other monitors throughout the network. The enterprise itself need not be stable in its configuration or centrally administered. Rather, it may exist as an emergent entity through the interconnections of the domains. EMERALD's ability to perform interdomain event analysis is vital to addressing more global, information warfare-like attacks against the entire enterprise (see Section IV).

A. The EMERALD Monitor

The generic EMERALD monitor architecture is illustrated in Figure 1. The architecture is designed to enable the flexible introduction and deletion of analysis engines from the monitor boundary as necessary. In its dual-analysis configuration, an EMERALD monitor instantiation combines signature analysis with statistical profiling to provide complementary forms of analysis over the operation of network services and infrastructure. In general, a monitor may include additional analysis engines that may implement other forms of event analysis, or a monitor may consist of only a single resolver implementing a response policy based on intrusion summaries produced by other EMERALD monitors. Monitors also incorporate a versatile application programmers' interface that enhances their ability to

interoperate with the analysis target, and with other third-party intrusion-detection tools.

Underlying the deployment of an EMERALD monitor is the selection of a target-specific event stream. The event stream may be derived from a variety of sources including audit data, network datagrams, SNMP traffic, application logs, and analysis results from other intrusion-detection instrumentation. The event stream is parsed, filtered, and formatted by the target-specific event-collection methods provided within the resource object definition (see Section III-B). Event records are then forwarded to the monitor's analysis engine(s) for processing.

EMERALD's *profiler engine* performs statistical profile-based anomaly detection given a generalized event stream of an analysis target (Section III-C). EMERALD's *signature engine* requires minimal state-management and employs a rule-coding scheme that breaks from traditional expert-system techniques to provide a more focused and distributed signature-analysis model (Section III-D). Multiple analysis engines implementing different analysis methods may be employed to analyze a variety of event streams that pertain to the same analysis target. These analysis engines are intended to develop significantly lower volumes of abstract intrusion or suspicion reports. The profiler and signature engines receive large volumes of event logs specific to the analysis target, and produce smaller volumes of intrusion or suspicion reports that are then fed to their associated *resolver*.

EMERALD's resolver is the coordinator of analysis reports and the implementor of the "response policy" (Section III-E). A resolver may correlate analysis results produced externally by other analysis engines to which it subscribes, and it may be bound to one or more analysis engines within the monitor boundary. Because the volume of its input is much lower than the event-stream volumes processed by the analysis engines, the resolver is able to implement sophisticated management and control policies over the analysis engines. The resolver also provides the primary interface between its associated analysis engines, the analysis target, and other intrusion-detection modules. In general, monitors may exist with multiple analysis engines, and support the capability to interoperate with third-party analysis engines.

At the center of the EMERALD monitor is a structure called a *resource object*. The resource object is a pluggable library of target-specific configuration data and methods that allows the monitor code-base to remain independent from the analysis target to which it is deployed (Section III-B). Customizing and dynamically configuring an EMERALD monitor thus becomes

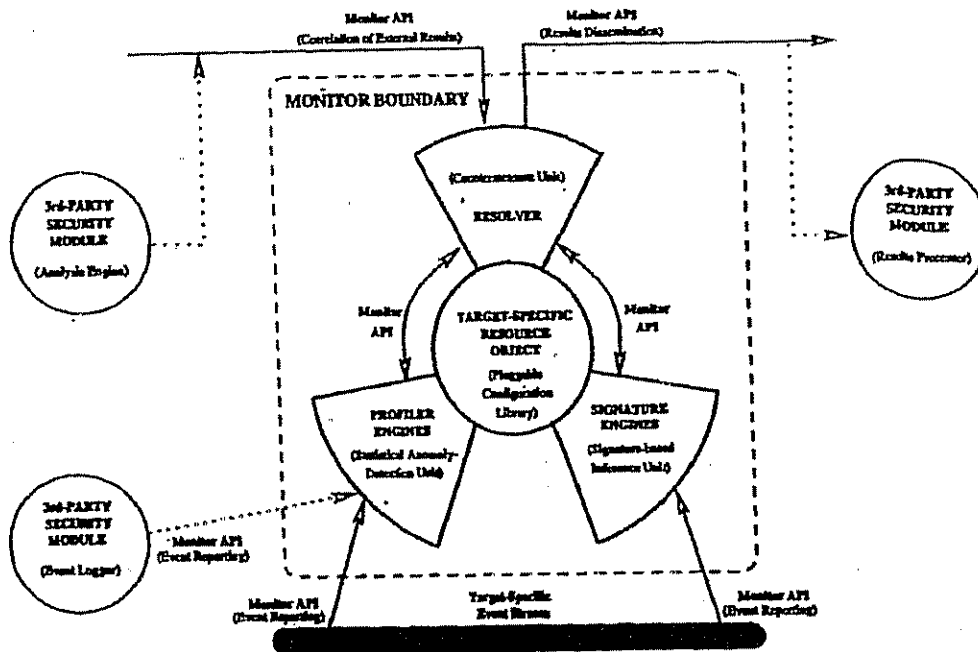


Fig. 1. The Generic EMERALD Monitor Architecture

a question of building and defining the fields of the analysis target's resource object.

Interoperability is especially critical to EMERALD's decentralized monitoring scheme, and extends within EMERALD's own architectural scope as well as to third-party modules. To support interoperability, EMERALD monitors incorporate a bidirectional messaging system. Section III-F discusses our efforts to develop a standard interface specification for communication within and between EMERALD monitors and external modules. Using this interface specification, third-party modules can communicate with EMERALD monitors in a variety of ways, as illustrated in Figure 1. Third-party modules operating as event-collection units may employ EMERALD's external interfaces to submit event data to the analysis engines for processing. Such third-party modules would effectively replace the monitor's own event-collection methods (Section III-B). Third-party modules may also submit and receive analysis results via the resolver's external interfaces. This will allow third-party modules to incorporate the results from EMERALD monitors into their own surveillance efforts, or to contribute their results to the EMERALD analysis hierarchy. Lastly, the monitor's internal API allows third-party analysis engines to be linked directly into the monitor boundary.

All EMERALD monitors (service, domain, and enterprise) are implemented using the same monitor code-base. The EMERALD monitor architecture is designed generally enough to be deployed at various abstract layers in the network. The only differences between deployed monitors are their resource object definitions. This reusable software architecture is a major project asset, providing significant benefits to the implementation and maintenance efforts. The following sections briefly describe the various components that make up the EMERALD monitor architecture.

B. Resource Objects: Abstracting Network Entities

Fundamental to EMERALD's design is the abstraction of the semantics of the analysis target from the EMERALD monitor. By logically decoupling the implementation of the EMERALD monitor from the analysis semantics of the analysis target, the extension of EMERALD's surveillance capabilities becomes a question of integration rather than implementation. The resource object contains all the operating parameters for each of the monitor's components as well as the analysis semantics (e.g., the profiler engine's measure and category definition, or the signature engine's penetration rule-base) necessary to process the target event stream. Once the resource object for a particular analysis target

is defined, it may be reused later by other EMERALD monitors that are deployed to equivalent analysis targets. For example, the resource object for a domain's router may be reused as other EMERALD monitors are deployed for other routers in the domain. A library of resource object definitions is being developed for commonly available network surveillance targets.

Figure 2 illustrates the general structure of the resource object. The resource object provides a pluggable configuration module for tuning the generic monitor code-base to a specific analysis target event stream. It minimally comprises the following variables (these variables may be extended as needed to accommodate the incorporation of new analysis engines into the monitor boundary):

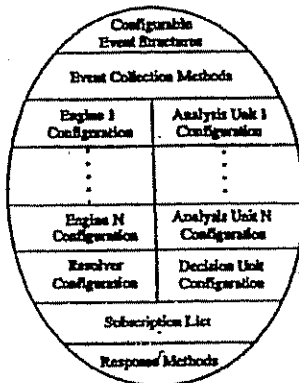


Fig. 2. The Generic EMERALD Monitor Architecture

- **Configurable Event Structures:** The monitor code-base maintains no internal dependence on the content or format of any given target event stream or the analysis results produced from analyzing the event stream. Rather, the resource object provides a universally applicable syntax for specifying the structure of event records and analysis results. Event records are defined based on the contents of the monitor's target event stream(s). Analysis result structures are used to package the findings produced by the analysis engine. Event records and analysis results are defined similarly to allow the eventual hierarchical processing of analysis results as event records by subscriber monitors.

- **Event-Collection Methods:** A set of filtering routines (or log conversion routines with custom filtering semantics) is employed by the analysis engines to gather and format target-specific event records. These are the native methods that interact directly with the system to parse the target event stream.

- **Engine N Configuration:** This refers to a col-

lection of variables and data structures that specifies the operating configuration of a fielded monitor's analysis engine(s). The resource object maintains a separate collection of operating parameters for each analysis engine instantiated within the monitor boundary.

- **Analysis Unit N Configuration:** Each analysis engine maintains an independently configured collection of intrusion-detection analysis procedures. This structure contains the configuration variables that define the semantics employed by the analysis engine to process the target-specific event stream.

- **Resolver Configuration:** The resource object maintains the operating parameters that specify the configuration of the resolver's internal modules.

- **Decision Unit Configuration:** This refers to the semantics used by the resolver's decision unit for merging the analysis results from the various analysis engines. The semantics include the response criteria used by the decision unit for invoking countermeasure handlers.

- **Subscription List:** This structure contains information necessary for establishing subscription-based communication sessions, which may include network address information and public keys used by the monitor to authenticate potential clients and servers. The subscription list field is an important facility for gaining visibility into malicious or anomalous activity outside the immediate environment of an EMERALD monitor. The most obvious examples where relationships are important involve interdependencies among network services that make local policy decisions. Consider, for example, the interdependencies between access checks performed during network file system mounting and the IP mapping of the DNS service. An unexpected mount monitored by the network file system service may be responded to differently if the DNS monitor informs the network file system monitor of suspicious updates to the mount requestor's DNS mapping.

- **Valid Response Methods:** Various response functions can be made available to the resolver as it receives intrusion reports from its analysis engines or intrusion summaries from subscribers. These are pre-programmed countermeasure methods that the resolver may invoke as intrusion summaries are received.

As discussed above, the fields of the resource object are defined and utilized during monitor initialization. In addition, these fields may be modified by internal monitor components, and by authorized external clients using the monitor's API. Once fields are modified, components can be requested to dynamically reload the configuration parameters defined in those fields. This gives EMERALD an important ability to provide adaptive

analysis a control functionality. However, it also introduces a potential stability problem if dynamic modifications are not tightly restricted to avoid cyclic modifications. To address this issue, monitors accept configuration requests from only immediate parents in EMERALD's analysis hierarchy.

C. Scalable Profile-Based Anomaly Detection

The original groundwork for SRI's IDES effort was performed over a decade ago. The first-generation statistics component was used to analyze System Management Facility (SMF) records from an IBM main-frame system [10] in the first half of the 1980s. Requirements for an anomaly-detection system that became IDES were documented in [6]. This research led to the development of the NIDES statistical profile-based anomaly-detection subsystem (NIDES/Stats), which employed a wide range of multivariate statistical measures to profile the behavior of individual users [9]. Analysis is user-based, where a statistical score is assigned to each user's session representing how closely currently observed usage corresponds to the established patterns of usage for that individual. The input source to the NIDES statistical component is an unfiltered and unsorted host audit log, which represents the activity of all users currently operating on the host.

In 1995, SRI conducted research under Trusted Information Systems' Safeguard project to extend NIDES/Stats to profile the behavior of individual applications [2]. Statistical measures were customized to measure and differentiate the proper operation of an application from operation that may indicate Trojan horse substitution. Under the Safeguard model, analysis is application-based, where a statistical score is assigned to the operation of applications and represents the degree to which current behavior of the application corresponds to its established patterns of operation. The Safeguard effort demonstrated the ability of statistical profiling tools to clearly differentiate the scope of execution among general-purpose applications. It also showed that statistical analyses can be very effective in analyzing activities other than individual users; by instead monitoring applications, the Safeguard analysis greatly reduced the required number of profiles and computational requirements, and also dramatically decreased the typical false-positive and false-negative ratios.

While NIDES/Stats has been reasonably successful profiling users and later applications, it will be extended to the more general subject class typography required by EMERALD. Nonetheless, the underlying mechanisms are well suited to the problem of network anomaly detection; with some adaptation. The required modifications center around extensive reworking

of NIDES/Stats to abstract and generalize its definition of measures and profiles, the streamlining of its profile management, and the adaptation of the configuration and reporting mechanisms to EMERALD's highly interoperable and dynamic message system interface.

The EMERALD profiler engine achieves total separation between profile management and the mathematical algorithms used to assess the anomaly of events. Profiles are provided to the computational engine as classes defined in the resource object. The mathematical functions for anomaly scoring, profile maintenance, and updating function in a fully general manner, not requiring any underlying knowledge of the data being analyzed beyond what is encoded in the profile class. The event-collection interoperability supports translation of elementary data (the analysis target's event stream) to the profile and measure classes. At that point, analysis for different types of monitored entities is mathematically similar. This approach imparts great flexibility to the analysis in that fading memory constants, update frequency, measure type, and so on are tailored to the entity being monitored.

Each profiler engine is dedicated to a specific target event stream at the elementary level. Such localized, target-specific analyses (unlike the monolithic approach employed by NIDES/Stats) provide a more distributed, building-block approach to monitoring, and allow profiling computations to be efficiently dispersed throughout the network. Because the event stream submitted to the profiler engine is specific to the analysis target's activity, profile management is greatly simplified, in that there is no need to support multisubject profile instantiations.

In addition, the results of service-layer profiler engines can be propagated to other monitors operating higher in EMERALD's layered analysis scheme, offering domain- or enterprise-wide statistical profiling of anomaly reports. Profiler engines may operate throughout the analysis hierarchy, further correlating and merging service-layer profiles to identify more widespread anomalous activity. The underlying mathematics are the same for each instance, and all required information specific to the entity being monitored (be it a network resource or other EMERALD monitors producing analysis results at lower layers in the analysis hierarchy) is entirely encapsulated in the objects of the profile class.

D. Scalable Signature Analysis

Signature analysis is a process whereby an event stream is mapped against abstract representations of event sequences that are known to indicate undesirable activity. However, simplistic event binding alone may not necessarily provide enough indication to ensure the

accurate detection of the target activity. Signature analyses must also distinguish whether an event sequence being witnessed is actually transitioning the system into the anticipated compromised state. Additionally, determining whether a given event sequence is indicative of an attack may be a function of the preconditions under which the event sequence is performed. To enable this finer granularity of signature recognition, previous efforts have employed various degrees of state detection and management logic (one such example is found in [18]). However, as discussed in Section II, the incorporation of sophisticated rule- and state-management features must be balanced with the need to ensure an acceptable level of performance.

In many respects, EMERALD's signature-analysis strategy departs from previous centralized rule-based efforts. EMERALD employs a highly distributed analysis strategy that, with respect to signature analysis, effectively modularizes and distributes the rule-base and inference engine into smaller, more focused signature engines. This has several benefits beyond the performance advantages from evenly distributing the computational load across network resources.

By narrowing the scope of activity in the event stream to a single analysis target, the noise ratio from event records that the signature engine must filter out is greatly reduced. This noise filtering of the event stream helps the signature engine avoid misguided searches along incorrect signature paths. EMERALD also partitions and distributes the signature activity representations. Rather than maintaining a central knowledge-base containing representations of all known malicious activity across a given computing environment, EMERALD distributes a tailored set of signature activity with each monitor's resource object.

EMERALD's signature-analysis objectives depend on which layer in EMERALD's hierarchical analysis scheme the signature engine operates. Service-layer signature engines attempt to monitor network services and infrastructure for attempts to subvert or misuse these components to penetrate or interfere with the domain's operation. Service monitors target external and perhaps unauthenticated individuals who attempt to subvert services or domain components to perform actions outside their normal operating scope. The EMERALD signature engine scans the event stream for events that represent attempted exploitations of known attacks against the service, or other activity that stands alone as warranting a response from the EMERALD monitor.

Above the service layer, signature engines scan the aggregate of intrusion reports from service monitors in an attempt to detect more global coordinated attack scenarios or scenarios that exploit interdependencies

among network services. The DNS/NFS attack discussed in Section III-B is one such example of an aggregate attack scenario. The fault-propagation model presented in [20] offers a general example of modeling interdependency of network assets (in this case fault interdependencies in a nonmalicious environment) that is also of general relevance for EMERALD's domain- and enterprise-layer intrusion correlation.

E. A Universal Resolver: Correlation and Response

EMERALD maintains a well-defined separation between analysis activities and response logic. Implementation of the response policy, including coordinating the dissemination of the analysis results, is the responsibility of the EMERALD resolver. The resolver is an expert system that receives the intrusion and suspicion reports produced by the profiler and signature engines, and based on these reports invokes the various response handlers defined within the resource object. Because the volume of intrusion and suspicion reports is lower than the individual event reports received by the analysis engines, the resolver can afford the more sophisticated demands of maintaining the configuration, and managing the response handling and external interfaces necessary for monitor operation. Furthermore, the resolver adds to the extensibility of EMERALD by providing the subscription interface through which third-party analysis tools can interact and participate in EMERALD's layered analysis scheme.

Upon its initialization, the resolver references various fields within the associated resource object. The resolver initiates authentication and subscription sessions with those EMERALD monitors whose identities appear in the resource object's subscription-list field. It also handles all incoming requests by subscribers, which must authenticate themselves to the resolver. (Details of EMERALD's subscription-session authentication process are discussed in [19].) Once a subscription session is established with a subscriber monitor, the resolver acts as the primary interface through which configuration requests are received, probes are handled, and intrusion reports are disseminated.

EMERALD supports extensive intermonitor sharing of analysis results throughout its layered analysis architecture. Resolvers are able to request and receive intrusion reports from other resolvers at lower layers in the analysis hierarchy. As analysis results are received from subscribers, they are forwarded via the monitor's event filters to the analysis engines. This tiered collection and correlation of analysis results allows EMERALD monitors to represent and profile more global malicious or anomalous activity that is not visible from the local monitoring of individual network services and assets

(see Section IV).

In addition to its external-interface responsibilities, the resolver operates as a fully functional decision engine, capable of invoking real-time countermeasures in response to malicious or anomalous activity reports produced by the analysis engines. Countermeasures are defined in the response-methods field of the resource object. Included with each valid response method are evaluation metrics for determining the circumstances under which the method should be dispatched. These response criteria involve two evaluation metrics: a threshold metric that corresponds to the measure values and scores produced by the profiler engine, and severity metrics correspond to subsets of the associated attack sequences defined within the resource object. The resolver combines the metrics to formulate its monitor's response policy. Aggressive responses may include direct countermeasures such as closing connections or terminating processes. More passive responses may include the dispatching of integrity-checking handlers to verify the operating state of the analysis target.

The resolver operates as the center of intramonitor communication. As the analysis engines build intrusion and suspicion reports, they propagate these reports to the resolver for further correlation, response, and dissemination to other EMERALD monitors. The resolver can also submit runtime configuration requests to the analysis engines, possibly to increase or decrease the scope of analyses (e.g., enable or disable additional signature rules) based on various operating metrics. These configuration requests could be made as a result of encountering other intrusion reports from other subscribers. For example, an intrusion report produced by a service monitor in one domain could be propagated to an enterprise monitor, which in turn sensitizes service monitors in other domains to the same activity.

Lastly, a critical function of the EMERALD resolver is to operate as the interface mechanism between the monitor administrator and the monitor itself. From the perspective of an EMERALD resolver, the administrator interface is simply a subscribing service to which the resolver may submit its intrusion summaries and receive probes and configuration requests. The administrative interface tool can dynamically subscribe and unsubscribe to any of the deployed EMERALD resolvers, as well as submit configuration requests and asynchronous probes as desired.

F. The EMERALD Message System

Interoperability is especially critical to the EMERALD design, which from conception promotes dynamic extensibility through a building-block approach to scal-

able network surveillance. EMERALD monitors incorporate a duplex messaging system that allows them to correlate activity summaries and countermeasure information in a distributed hierarchical analysis framework. EMERALD's messaging system must address interoperability both within its own architectural scope and with other third-party analysis tools. To do this, the messaging system provides a well-defined programmer's interface that supports the bidirectional exchange of analysis results and configuration requests with alternative security tools.

EMERALD's message system operates under an asynchronous communication model for handling results dissemination and processing that is generically referred to as subscription-based message passing.² EMERALD component interoperation is client/server-based, where a client module may subscribe to receive event data or analysis results from servers. Once the subscription request is accepted by the server, the server module forwards events or analysis results to the client automatically as data becomes available, and may dynamically reconfigure itself as requested by the client's control requests. While this asynchronous model does not escape the overhead needed to ensure reliable delivery, it does reduce the need for client probes and acknowledgments.

An important goal in the design of EMERALD's interface specification is that the interface remain as implementation neutral as possible. To support an implementation-neutral communication framework, the message system is designed with strong separation between the programmer's interface specification and the issues of message transport.³ The interface specification embodies no assumptions about the target intrusion-detection modules, implementation languages, host platform, or network. The transport layer is architecturally isolated from the internals of EMERALD monitors so that transport modules may be readily introduced and replaced as protocols and security requirements are negotiated between module developers. The following briefly summarizes EMERALD's interface specification and transport layer design.

Interface Specification: Interface specification involves the definition of the messages that the various intrusion-detection modules must convey to one another, and how these messages should be processed. The message structure and content are specified in a completely implementation-neutral context. Internally, EMERALD monitors contain three general module types: event collection methods that collect and fil-

² Other communities have employed subscription-based push/pull data flow schemes for information such as network management traffic and WWW content.

³ Details of EMERALD's programmer's interface specification and transport design are provided in [19].

ter the target event stream, analysis engines that process the filtered events, and a resolver that processes and responds to the analysis engine results. Externally, EMERALD monitors interoperate with one another in a manner analogous to internal communication: service monitors produce local analysis results that are passed to the domain monitor; domain monitors correlate service monitor results, producing new results that are further propagated to enterprise monitors; enterprise monitors correlate and respond to the analysis results produced by domain monitors.

Both intramonitor and intermonitor communication employ identical subscription-based client-server models. With respect to intermonitor communication, the resolver operates as a client to the analysis engines, and the analysis engines operate as clients to the event filters. Through the internal message system, the resolver submits configuration requests and probes to the analysis engines, and receives from the analysis engines their analysis results. The analysis engines operate as servers providing the resolver with intrusion or suspicion reports either asynchronously or upon request. Similarly, the analysis engines are responsible for establishing and maintaining a communication link with a target event collection method (or event filter) and prompting the reconfiguration of the collection method's filtering semantics when necessary. Event collection methods provide analysis engines with target-specific event records upon which the statistical and signature analyses are performed.

Intermonitor communication also operates using the subscription-based hierarchy. A domain monitor subscribes to the analysis results produced by service monitors, and then propagates its own analytical results to its parent enterprise monitor. The enterprise monitor operates as a client to one or more domain monitors, allowing them to correlate and model enterprise-wide activity from the domain-layer results. Domain monitors operate as servers to the enterprise monitors, and as clients to the service-layer monitors deployed throughout their local domain. This message scheme would operate identically if correlation were to continue at higher layers of abstraction beyond enterprise analysis.

EMERALD's intramonitor and intermonitor programming interfaces are identical. These interfaces are subdivided into five categories of interoperation: channel initialization and termination, channel synchronization, dynamic configuration, server probing, and report/event dissemination. Clients are responsible for initiating and terminating channel sessions with servers.

Furthermore, clients are responsible for managing channel synchronization in the event of errors in message sequencing or periods of failed or slow response (i.e.,

"I'm alive" confirmations). Clients may also submit dynamic configuration requests to servers. For example, an analysis engine may request an event collection method to modify its filtering semantics. Clients may also probe servers for report summaries or additional event information. Lastly, servers may send clients intrusion/suspicion summaries or event data in response to client probes or in an asynchronous dissemination mode.

Transport Layer: The second part of the message system framework involves the specification of the transport mechanism used to establish a given communication channel between monitors or possibly between a monitor and a third-party security module. All implementation dependencies within the message system framework are addressed by the pluggable transport modules. Transport modules are specific to the participating intrusion-detection modules, their respective hosts, and potentially to the network—should the modules require cross-platform interoperation. Part of the integration of a monitor into a new analysis target is the incorporation of the necessary transport module(s) (for both internal and external communication).

It is at the transport layer where EMERALD addresses issues of communications security, integrity, and reliability. While it is important to facilitate interoperability among security mechanisms, this interoperability must be balanced with the need to ensure an overall level of operational integrity, reliability, and privacy. An essential element in the EMERALD messaging system design is the integration of secure transport to ensure a degree of internal security between EMERALD components and other cooperative analysis units.

The transport modules that handle intramonitor communication may be different from the transport modules that handle intermonitor communication. This allows the intramonitor transport modules to address security and reliability issues differently than how the intermonitor transport modules address security and reliability. While intramonitor communication may more commonly involve interprocess communication within a single host, intermonitor communication will most commonly involve cross-platform networked interoperation. For example, the intramonitor transport mechanisms may employ unnamed pipes [14], which provides a kernel-enforced private interprocess communication channel between the monitor components (this assumes a process hierarchy within the monitor architecture). The monitor's external transport, however, will more likely export data through untrusted network connections and thus require more extensive security management. To ensure the security and integrity of the message exchange, the external transport may employ

public/private key authentication protocols and session key exchange. Using this same interface, third-party analysis tools may authenticate and exchange analysis results and configuration information with EMERALD monitors in a well-defined, secure manner.

The pluggable transport allows EMERALD flexibility in negotiating security features and protocol usage with third parties. Of particular interest to the monitoring of network events is our planned incorporation of a commercially available network management system as a third-party module. That system will deliver monitoring results relating to security, reliability, availability, performance, and other attributes. The network management system may in turn subscribe to EMERALD results in order to influence network reconfiguration. This experiment will demonstrate the interoperation of intrusion-detection instrumentation with analysis tools that themselves do not specifically address security management.

IV. EMERALD NETWORK DEPLOYMENT

The EMERALD reusable-monitor architecture provides a framework for the organization and coordination of distributed event analysis across multiple administrative domains. EMERALD introduces a service-oriented, layered approach to representing, analyzing, and responding to network misuse. EMERALD's profiling and signature analyses are not performed as monolithic analyses over an entire domain, but rather are deployed sparingly throughout a large enterprise to provide focused protection of key network assets vulnerable to attack. This model leads to greater flexibility whenever the network configuration changes dynamically, and to improved performance, where computational load is distributed efficiently among network resources.

Domains under EMERALD surveillance are able to detect malicious activity targeted against their network services and infrastructure, and disseminate this information in a coordinated and secure way to other EMERALD monitors (as well as third-party analysis tools) distributed throughout the network. Reports of problems found in one domain can propagate to other monitors throughout the network using the subscription process. EMERALD's subscription-based communication strategy provides mutual authentication between participants, as well as confidentiality and integrity for all intermonitor message traffic (see Section III-F).

EMERALD's analysis scheme is highly composable, beginning at the service layer where EMERALD monitors analyze the security-relevant activity associated with an individual network service or network infrastructure. As service-layer monitors detect activity that

indicates possible misuse, this information is responded to by the monitor's local resolver to ensure immediate response. Misuse reports are also disseminated throughout EMERALD's web of surveillance, to the monitor's pool of subscribers.

Domain-layer monitors model and profile domain-wide vulnerabilities not detectable from the narrow visibility of the service layer. Domain monitors search for intrusive and anomalous activity across a group of interdependent service-layer components, subscribing to each service's associated service monitor. Domain monitors also operate as the dissemination point between the domain's surveillance and the external network surveillance. Where mutual trust among domains exists, domain monitors may establish peer relationships with one another. Peer-to-peer subscription allows domain monitors to share intrusion summaries from events that have occurred in other domains. Domain monitors may use such reports to dynamically sensitize their local service monitors to malicious activity found to be occurring outside the domain's visibility. Domain monitors may also operate within an enterprise hierarchy, where they disseminate intrusion reports to enterprise monitors for global correlation. Where trust exists between domains, peer-to-peer subscription provides a useful technique for keeping domains sensitized to malicious activity occurring outside their view.

Enterprise-layer monitors attempt to model and detect coordinated efforts to infiltrate domain perimeters or prevent interconnectivity between domains. Enterprise surveillance may be used where domains are interconnected under the control of a single organization, such as a large privately owned WAN. Enterprise surveillance is very similar to domain surveillance: the *enterprise monitor* subscribes to various domain monitors, just as the domain monitors subscribed to various local service monitors. The enterprise monitor (or monitors, as it would be important to avoid centralizing any analysis) focuses on network-wide threats such as Internet worm-like attacks, attacks repeated against common network services across domains, or coordinated attacks from multiple domains against a single domain. As an enterprise monitor recognizes commonalities in intrusion reports across domains (e.g., the spreading of a worm or a mail system attack repeated throughout the enterprise), its resolver can take steps to help domains counter the attack, and can also help sensitize other domains to such attacks before they are affected.

EMERALD's distributed analysis paradigm provides several significant performance advantages over the centralized signature analysis and statistical profiling tools from which its architecture is derived. In a large network, event activity is dispersed throughout its spa-

tially distributed components, occurring in parallel and in volumes that are difficult for centralized analysis tools to manage. EMERALD distributes the computational load and space utilization needed to monitor the various network components, and performs its analysis and response activity locally. Local detection and response also helps to ensure timely protection of network assets. Furthermore, EMERALD's distributed monitor deployment effectively parallelizes the statistical profiling and signature analyses. Once the event streams from the various analysis targets are separated and submitted to the deployed monitors, event correlation, profiling, and response handling are all managed by independent computational units. Lastly, EMERALD's dynamic extensibility allows an integrator to selectively choose the key elements in a network that require monitoring, and the ability to alter analysis coverage dynamically.

V. RELATED WORK

EMERALD is not intended as a replacement to more centralized, host-based, user-oriented intrusion-detection tools, but rather as a complementary architecture that addresses threats from the interconnectivity of domains in hostile environments. Specifically, EMERALD attempts to detect and respond to both anticipated and unanticipated misuses of services and infrastructure in large network-based enterprises, including external threats that attempt to subvert or bypass a domain's network interfaces and control mechanisms to gain unauthorized access to domain resources or prevent the availability of these resources. EMERALD also provides a framework for recognizing more global threats to interdomain connectivity, including coordinated attempts to infiltrate or destroy connectivity across an entire network enterprise. A more detailed discussion of EMERALD's relationship with other work is given in [19]. Here, we merely allude to its position in the spectrum of research in intrusion detection, fault detection, and alarm correlation.

A. Related Intrusion Detection Research

EMERALD considerably generalizes and extends the earlier pioneering work of SRI's IDES and NIDES [1], overcoming previous limitations with respect to scalability, applicability to networking, interoperability, and inability to detect distributed coordinated attacks. It generalizes to network environments the Safeguard experience [2], which overcame profile explosion and scalability problems by locally profiling the activities of subsystems and commands rather than of individual users. EMERALD also extends the statistical-profile model of NIDES, to analyze the operation of network services, network infrastructure, and activity reports from other

EMERALD monitors. Various other efforts have considered one of the two types of analysis – signature-based (e.g., Porras [18] has used a state-transition approach; the U.C. Davis and Trident DIDS [4] addresses abstracted analysis for networking, but not scalability; the Network Security Monitor [7] seeks to analyze packet data rather than conventional audit trails; Purdue [5] seeks to use adaptive-agent technology) or profile-based. More recent work in UC Davis' GrIDS effort [24] employs *activity graphs* of network operations to search for traffic patterns that may indicate network-wide coordinated attacks. (Ko has considered writing specifications for expected behavior [13], which is sort of a compromise between signature analysis and behavioral profiling.)

B. Related Research in Fault Detection

EMERALD is somewhat similar conceptually to various efforts in alarm correlation and high-volume event correlation/fault detection in the network management community [8], [15], [16]. EMERALD's architecture and layered analysis is somewhat similar to the distributed event correlation system (DECS) discussed in [12]. However, DECS makes several simplifications in its stateless event modeling scheme that do not translate well to a malicious environment for detecting intrusions. Recent work in nonmalicious fault isolation [20] is also relevant, and is being considered. However, none of these efforts shames EMERALD's abilities for recursive hierarchical abstraction and misuse detection, nor do they include provisions to ensure their own survivability in hostile environments.

VI. CONCLUSIONS

This paper introduces EMERALD, a composable surveillance and response architecture oriented toward the monitoring of distributed network elements. EMERALD targets external threat agents who attempt to subvert or bypass network interfaces and controls to gain unauthorized access to domain resources. EMERALD builds a multiple local monitoring capability into a framework for coordinating the dissemination of distributed analyses to provide global detection and response to network-wide coordinated attacks. The basic analysis unit in this architecture is the EMERALD monitor, which incorporates both signature analysis and statistical profiling. By separating the analysis semantics from the analysis and response logic, EMERALD monitors can be easily integrated throughout EMERALD's layered network surveillance strategy.

EMERALD builds on and considerably extends past research and development in anomaly and misuse detection, to accommodate the monitoring of large dis-

tributed systems and networks. Because the real-time analysis itself can be distributed and applied where most effective at different layers of abstraction, EMERALD has significant advantages over more centralized approaches in terms of event detectability and response capabilities, and yet can be computationally realistic. It can detect not only local attacks, but also coordinated attacks such as distributed denials of service. The EMERALD design addresses interoperability within its own scope, and in so doing enables its interoperability with other analysis platforms as well. EMERALD's inherent generality and flexibility in terms of what is being monitored and how the analysis is accomplished suggests that the design can be readily extended to monitoring other attributes such as survivability, fault tolerance, and assured availability.

REFERENCES

- [1] D. Anderson, T. Frivold, and A. Valdes. Next-generation intrusion-detection expert system (NIDES). Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA, SRI-CSL-95-07, May 1995.
- [2] D. Anderson, T. Lunt, H. Javitz, A. Tamaru, and A. Valdes. Safeguard final report: Detecting unusual program behavior using the NIDES statistical component. Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA, 2 December 1993.
- [3] J.P. Anderson. Computer security technology planning study. Technical Report ESD-TR-73-51, ESD/AFSC, Hanscom AFB, Bedford, MA, October 1972.
- [4] J. Brentano, S.R. Snapp, G.V. Dias, T.L. Goan, L.T. Heberlein, C.H. Ho, K.N. Levitt, and B. Mukherjee. An architecture for a distributed intrusion detection system. In *Fourteenth Department of Energy Computer Security Group Conference*, pages 25-45 in section 17, Concord, CA, May 1991. Department of Energy.
- [5] M. Crosbie and E.H. Spafford. Active defense of a computer system using autonomous agents. Technical report, Department of Computer Sciences, CSD-TR-95-008, Purdue University, West Lafayette IN, 1995.
- [6] D.E. Denning and P.G. Neumann. Requirements and model for IDES - a real-time intrusion-detection expert system. Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA, August 1985.
- [7] T.L. Heberlein, B. Mukherjee, and K.N. Levitt. A method to detect intrusive activity in a networked environment. In *Proceedings of the Fourteenth National Computer Security Conference*, pages 362-371, Washington, D.C., 1-4 October 1991. NIST/NCSC.
- [8] G. Jakobson and M.D. Weissman. Alarm correlation. *IEEE Network*, pages 52-59, November 1993.
- [9] H.S. Javitz and A. Valdes. The NIDES statistical component description and justification. Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA, March 1994.
- [10] H.S. Javitz, A. Valdes, D.E. Denning, and P.G. Neumann. Analytical techniques development for a statistical intrusion-detection system (SIDS) based on accounting records. Technical report, SRI International, Menlo Park, CA, July 1988.
- [11] P.M. Joyal. Industrial espionage today and information wars of tomorrow. In *National Information Systems Security Conference*, Baltimore, Maryland, pages 139-150, Washington, D.C., 22-25 October 1996.
- [12] S. Kliger, S. Yemini, Y. Yemini, D. Ohsie, and S. Stolfo. A coding approach to event correlation. In *Proceedings of the Fourth International Symposium on Integrated Network Management (IFIP/IEEE)*, Santa Barbara, CA, May 1995, pages 266-277. Chapman & Hall, London, England, 1995.
- [13] C. Ko, M. Ruschitzka, and K. Levitt. Execution monitoring of security-critical programs in distributed systems: A specification-based approach. In *Proceedings of the 1997 Symposium on Security and Privacy*, pages 175-187, Oakland, CA, May 1997. IEEE Computer Society.
- [14] D. Levine. *POSIX Programmer's Guide*. O'Reilly and Associates Incorporated, 1991.
- [15] M. Mansouri-Samani and M. Sloman. Monitoring distributed systems. *IEEE Network*, pages 20-30, November 1993.
- [16] K. Meyer, M. Erlinger, J. Belsaer, C. Sunahine, G. Goldszmidt, and Y. Yemini. Decentralizing control and intelligence in network management. In *Proceedings of the Fourth International Symposium on Integrated Network Management (IFIP/IEEE)*, Santa Barbara, CA, May 1995, pages 4-16. Chapman & Hall, London, England, 1995.
- [17] P.G. Neumann. *Computer-Related Risks*. ACM Press, New York, and Addison-Wesley, Reading, MA, 1994. ISBN 0-201-55805-X.
- [18] P.A. Porras and R.A. Kemmerer. Penetration state transition analysis: A rule-based intrusion detection approach. In *Proceedings of the Eighth Annual Computer Security Applications Conference (San Antonio, TX, Nov.30-Dec.4)*, pages 220-229. IEEE, 1992.
- [19] P.A. Porras and P.G. Neumann. Conceptual design and planning for EMERALD: event monitoring enabling responses to anomalous live disturbances. Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA, October 1997. Available for download via <http://www.csl.sri.com/intrusion.html>.
- [20] L. Ricculli and N. Shacham. Modelling correlated alarms in network management systems. *Communications Networks and Distributed Systems Modeling and Simulation*, 1997.
- [21] J.A. Rochlis and M.W. Eichin. With microscope and tweezers: The Worm from MIT's perspective. *Communications of the ACM*, 32(6):689-693, June 1989.
- [22] B. Rosen. Vulnerabilities of network control protocols. *ACM SIGSOFT Software Engineering Notes*, 8(1):6-8, January 1981.
- [23] E.H. Spafford. The Internet Worm: crisis and aftermath. *Communications of the ACM*, 32(6):678-687, June 1989.
- [24] S. Staniford-Chen, S. Cheung, R. Crawford, J. Frank M. Dillger, J. Hoagland, K. Levitt, G. Wee, R. Yip, and D. Zerkel. Grids—a graph based intrusion detection system for large networks. In *Proceedings of the Nineteenth National Information Systems Security Conference*, pages 361-370, October 1996.

EXHIBIT D

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

SRI INTERNATIONAL, INC., a California
Corporation,

Plaintiff,

v.

INTERNET SECURITY SYSTEMS, INC.,
a Delaware corporation, INTERNET
SECURITY SYSTEMS, INC., a Georgia
corporation, and SYMANTEC
CORPORATION, a Delaware corporation,

Defendants.

Case No. 04-1199-SLR

**PLAINTIFF SRI INTERNATIONAL, INC.'S FIRST SET OF REQUESTS FOR
PRODUCTION TO DEFENDANT INTERNET SECURITY SYSTEMS, INC., A
GEORGIA CORPORATION [NOS. 1 – 81]**

Pursuant to Rules 26 and 34 of the Federal Rules of Civil Procedure, Plaintiff SRI International, Inc. hereby requests Defendant Internet Security Corporation, a Georgia corporation ("ISS-GA"), to respond to these Requests by producing for inspection and copying the Documents and Things requested below at the offices of Fish & Richardson P.C., 500 Arguello Street, Suite 500, Redwood City, California 94063, within thirty (30) days of service of these Requests, or at such other times and places as counsel for the parties may agree.

INSTRUCTIONS

1. When producing a Document, please produce the Document as it is kept in the ordinary course of business, or indicate the paragraph of these Requests to which that Document is responsive.

2. Electronic records and computerized information must be produced in an intelligible format or together with a description of the system from which it was derived sufficient to permit rendering the materials intelligible.

3. In producing Documents, furnish all Documents known or available to ISS-GA regardless of whether such Documents are possessed directly by ISS-GA, or by any parent, subsidiary, or affiliated corporation, or any of ISS-GA's officers, directors, employees, agents, representatives, or attorneys, as well as any other Documents in ISS-GA's custody or control.

4. File folders with tabs or labels identifying Documents called for by these Requests must be produced intact with such Documents.

5. Selection of Documents from the files and other sources and the numbering of such Documents shall be performed in such a manner as to insure that the source of each Document may be determined, if necessary.

6. Documents attached to each other must not be separated.

7. The term "all Documents" means any and all Documents that might reasonably be located through a search of all locations reasonably likely to contain Documents called for by these Requests.

8. Should any Document be withheld based on some limitation of discovery (Including a claim of privilege), please supply the following information:

(a) The identity(ies) of each Document's author(s), writer(s), sender(s), or initiator(s);

(b) The identity(ies) of each Document's recipient(s), addressee(s), or party(ies) for whom it was intended;

(c) The date of creation or transmittal indicated on each Document, or an estimate of that date, indicated as such, if no date appears on the Document;

(d) The general subject matter as described on each Document, or, if no such description appears, then some other description sufficient to identify the Document; and

(e) The claimed ground(s) for limitation of discovery (e.g., "attorney-client privilege" or "attorney work product doctrine").

9. The written answer to each individual request for production must repeat verbatim, immediately before each answer, the text of the individual request for production being answered.

DEFINITIONS

As used in these Requests, the following terms have the meanings indicated:

1. "SRP" means SRI International, Inc., including its officers, directors, employees, agents, and attorneys.

2. "ISS-GA," "Defendant," "you," or "your" means Internet Security Systems, a Georgia corporation, including its past and present officers, directors, employees, consultants, agents, and attorneys and others acting or purporting to act on its behalf, and including their predecessors, subsidiaries, parents, and affiliates.

3. The phrase the "'615 patent" means U.S. Patent No. 6,711,615.

4. The phrase the "'203 patent" means U.S. Patent No. 6,484,203.

5. The phrase the "'338 patent" means U.S. Patent No. 6,321,338.

6. The phrase the "'212 patent" means U.S. Patent No. 6,708,212.

7. The phrase the "'874 patent" means U.S. Patent No. 6,704,874.

8. The phrase "Patents-in-Suit" means the '615, '203, '338, '212, and '874 Patents.

9. The word "Document" is used herein in its broadest sense to include everything that is contemplated by Rule 26 and Rule 34 of the Federal Rules of Civil Procedure, Including Documents stored in hard copy or electronic form. Electronic Documents Include electronic mail, computer source code, object code, and microcode, and Documents stored on any media accessible by electronic means. A comment or notation appearing on any "Document" and not a part of the original text is to be considered a separate "Document."

10. "Thing" means any tangible object other than a Document.

11. "Person" or "Persons" include not only natural individuals, but also, without limitation, firms, partnerships, associations, corporations, and other legal entities, and divisions, departments, or other units thereof.

12. These "Requests" shall mean the instant document, i.e., Plaintiff SRI International, Inc.'s First Set of Requests for Production of Documents and Things to Defendant Internet Security Systems, A Georgia Corporation [Nos. 1 – 81].

13. "Including" shall mean "including but not limited to."

14. The terms "and" and "or" shall be construed conjunctively or disjunctively, whichever makes the individual request more inclusive.

15. The singular and masculine form of any noun or pronoun shall embrace and be read and applied as embracing the plural, the feminine, and the neuter, except where circumstances clearly make it inappropriate.

16. The terms "refer," "referring," "relate," or "relating" as used herein include, but are not limited to the following meanings: bearing upon, concerning, constituting, discussing, describing, evidencing, identifying, concerning, mentioning, in connection with, pertaining to, respecting, regarding, responding to, or in any way factually or logically relevant to the matter described in the request.

17. The term "Accused Products" shall mean all software products relating to computer network intrusion detection, prevention or analysis developed, made, used, sold or offered for sale in, or imported into the United States at any time from 2001 to the present, including, but not limited to, SiteProtector and Proventia.

DOCUMENTS AND THINGS TO BE PRODUCED

REQUEST FOR PRODUCTION NO. 1:

All documents referring or relating in any way to any of your document, file, or record retention or destruction policies, including without limitation, any such policies referring or relating to electronic data, including source code, from 2001 to the present.

REQUEST FOR PRODUCTION NO. 2:

Management and/or organization charts setting forth the entities and/or the names and/or titles of individuals involved in the decision to develop, development, manufacture, marketing, sale and distribution of the Accused Products.

REQUEST FOR PRODUCTION NO. 3:

All brochures, advertisements, press releases, price lists, catalogs and other marketing and promotional materials relating to the Accused Products.

REQUEST FOR PRODUCTION NO. 4:

All documents containing, describing or relating to any communication between you and any third-party supplier of pieces of software, including source code, or other technology incorporated in the Accused Products.

REQUEST FOR PRODUCTION NO. 5:

All documents that describe, illustrate, or depict names and/or functions of subsidiaries, departments or affiliated entities related to you that were involved in any matter in the decision to develop, conception, design, manufacture, testing, marketing or sale of any of the Accused Products.

REQUEST FOR PRODUCTION NO. 6:

All documents that refer or relate to the research, development, and testing of the Accused Products, including but not limited to drawings, prototypes, notes, notebooks, workbooks, project reports, correspondence, memoranda, test results, schematics, flow charts and invention disclosures.

REQUEST FOR PRODUCTION NO. 7:

All operation manuals, user manuals, installation guides, dealers guides, and service manuals for the Accused Products.

REQUEST FOR PRODUCTION NO. 8:

All documents relating to the decision to develop and market the Accused Products, including all market studies, communication with customers, and internal correspondence.

REQUEST FOR PRODUCTION NO. 9:

All documents that refer or relate to the design of the Accused Products, including but not limited to specifications, data sheets, drawings, diagrams, schematics, manuals, notes, notebooks, workbooks, correspondence, and memoranda.

REQUEST FOR PRODUCTION NO. 10:

All documents sufficient to show the structure, function, or operation of the Accused Products, including but not limited to drawings, block diagrams and operating instructions.

REQUEST FOR PRODUCTION NO. 11:

All documents that refer or relate to the research, development, and testing of the method and/or process for using any Accused Product or products incorporating the same, including but not limited to notes, notebooks, workbooks, correspondence, memoranda, and test results.

REQUEST FOR PRODUCTION NO. 12:

All documents that refer or relate to the design of the method and/or process for using any Accused Product or products incorporating the same, including but not limited to instructions for use, drawings, diagrams, manuals, notes, notebooks, workbooks, correspondence, and memoranda.

REQUEST FOR PRODUCTION NO. 13:

All documents that refer or relate to the manufacture, , and production of the Accused Products, from 2001 to the present, including but not limited to the location of all manufacturing, and production facilities and the quantity of each Accused Product manufactured or produced at each such facility.

REQUEST FOR PRODUCTION NO. 14:

All documents, including but not limited to promotional materials, press releases, trade journal articles, advertisements, catalogs, documents prepared for trade shows and meetings, package inserts, technical data sheets, specifications, price lists, sales presentations, and sales and marketing forecasts and projections, that refer or relate to advertising and marketing of the Accused Products from 2001 to the present.

REQUEST FOR PRODUCTION NO. 15:

All documents that refer or relate to the sale within the United States of the Accused Products, from 2001 to the present, including but not limited to documents sufficient to show ISS-GA's actual and expected gross revenues, net profits, gross revenues, sales, sales prices, and returns of the Accused Products.

REQUEST FOR PRODUCTION NO. 16:

All documents that have been or are included with the sale of the Accused Products or products containing same, including without limitation instruction manuals, instructions for use, and users' guides.

REQUEST FOR PRODUCTION NO. 17:

All documents referring or relating to the distribution network maintained by ISS-GA for purposes of selling, distributing, or licensing the Accused Products and products containing same.

REQUEST FOR PRODUCTION NO. 18:

All documents that refer or relate to the importation into the United States of the Accused Products and any product incorporating an Accused Product, from 2001 to the present, including but not limited to the quantity of Accused Products (specifically identified by product) imported into the United States, the dates of such importation, and the ports of entry.

REQUEST FOR PRODUCTION NO. 19:

All documents that refer or relate to any United States inventory of the Accused Products, from 2001 to the present, including but not limited to the location of inventory, the amount of inventory, and the owner and/or holder of each such inventory.

REQUEST FOR PRODUCTION NO. 20:

All documents sufficient to show or relate to any step in the method and/or process of using the Accused Products and products incorporating the Accused Products, including but not limited to drawings, and block diagrams.

REQUEST FOR PRODUCTION NO. 21:

Documents sufficient to identify each and every Accused Product that has been imported into the United States, sold for importation into the United States, or sold or offered for sale after importation into the United States, from 2001 to the present.

REQUEST FOR PRODUCTION NO. 22:

Documents sufficient to identify each and every Accused Product that is being planned, that has been designed, and/or that is in development that is capable of being used in a product, circuit, or device that has been imported into the United States, sold for importation into the United States, or sold or offered for sale after importation into the United States.

REQUEST FOR PRODUCTION NO. 23:

Five samples of each version of each Accused Product, and all literature or documentation distributed with each Accused Product, any additional products that are required for the intended uses (both inside the United States and outside the United States) of each Accused Product, and all instruction or other guidance information associated with the sale, demonstration or use of each Accused Product.

REQUEST FOR PRODUCTION NO. 24:

All documents that refer or relate to instructions for using the Accused Products, including prior versions of instructions for use and instructions for use provided with the

sale, demonstration, marketing, or training for use of the Accused Products, both inside the United States and outside the United States.

REQUEST FOR PRODUCTION NO. 25:

All documents discussing or making reference to the quality, value, acceptability, workability, performance or benefits of any Accused Product, including communications with anyone who has used a Accused Product integrated circuit product praising, criticizing, or otherwise commenting on the Accused Products.

REQUEST FOR PRODUCTION NO. 26:

All documents that refer or relate to any failures, problems with and/or complaints about any of the Accused Products and products containing same or the method and/or process of using any Accused Product and products containing same.

REQUEST FOR PRODUCTION NO. 27:

All documents that refer or relate to uses, tests, or evaluations of the Accused Products and products containing same or the method and/or process of using any Accused Product and products containing same performed or conducted by ISS-GA, ISS-GA's customers, or any other third party that ISS-GA is aware of.

REQUEST FOR PRODUCTION NO. 28:

All documents that refer or relate to training any customer in the use of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 29:

All documents evidencing contributions made by any of your employees to the conception, design, or development of the Accused Products.

REQUEST FOR PRODUCTION NO. 30:

All United States and foreign patents and patent applications, including all documents referring or relating to such patents and patent applications, owned or controlled by you that relate to the Accused Products.

REQUEST FOR PRODUCTION NO. 31:

All documents constituting or relating to licenses or indemnification agreements between you and any other party relating to the Accused Products, or comparable products.

REQUEST FOR PRODUCTION NO. 32:

All articles, speeches, presentations or interviews that have been written or given by your employees, officers, directors or other representatives that refer or relate to the Accused Products.

REQUEST FOR PRODUCTION NO. 33:

All documents referring or relating to the Patents-in-Suit, including, but not limited to, documents that refer or relate to when and how you first became aware of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 34:

All documents found or identified during any enforceability searches, infringement analyses, prior art searches or studies related to the Patents-in-Suit, including any copies of patents, publications or other prior art.

REQUEST FOR PRODUCTION NO. 35:

All documents referring or relating to or constituting any patent or publication that you contend invalidates any of the claims of any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 36:

All documents, including but not limited to opinion letters, memoranda, or other documents that refer to or relate to the validity, infringement and/or enforceability of any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 37:

All internal communications among and between your personnel that reflect, refer to, or concern any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 38:

All documents, including without limitation articles and internal technical memoranda authored by you that relate to the subject matter described, disclosed or claimed in any of the patents-in-suit.

REQUEST FOR PRODUCTION NO. 39:

All communications or opinions of your officers, directors and/or employees regarding the infringement, validity or enforceability of any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 40:

All documents referring or relating to any consideration given by ISS-GA to any possible liability for patent infringement should ISS-GA commence or continue to make, use, distribute or sell the Accused Products or products containing same, including without limitation any internal notes or memoranda.

REQUEST FOR PRODUCTION NO. 41:

All documents that set forth on a monthly, quarterly or annual basis from 2001 to the present, the number and nature of units developed, used, sold and/or distributed for each of the Accused Products.

REQUEST FOR PRODUCTION NO. 42:

All documents that set forth, on a monthly, quarterly or annual basis from 2001 to present, the amount of sales in U.S. dollars or local currency (indicate currency type if not U.S. dollars) for each of the Accused Products.

REQUEST FOR PRODUCTION NO. 43:

All projections, forecasts, market share reports, marketing acceptance documents business plans, strategic plans, fiscal plans, marketing plans, or sales plans that refer or relate to the sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 44:

All documents that, on a monthly, quarterly or annual basis, refer or relate to unit sales, gross selling price, net selling price, or amounts deducted from gross selling price to arrive at net selling price, for each of the Accused Products.

REQUEST FOR PRODUCTION NO. 45:

All profit and loss statements, printouts, statements or other documents that set forth or include gross profit figures for the Accused Products on a monthly, quarterly or annual basis from 2001 to present.

REQUEST FOR PRODUCTION NO. 46:

All annual, semi-annual, quarterly, monthly or other documents that set forth or include your gross expenses incurred in the manufacture, distribution and sale of the Accused Products, including but not limited to, direct labor costs, direct manufacturing costs, selling costs, variable overhead costs, and all other costs associated with the manufacture, distribution and sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 47:

Documents that evidence the date of the first sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 48:

All documents that identify any of your divisions, affiliates, subsidiaries, or suppliers that have developed or manufactured any of the Accused Products.

REQUEST FOR PRODUCTION NO. 49:

Financial statements, including profit and loss statements, income statements, balance sheets, statements of cash flow, statements of retained earnings and notes thereto for any of your affiliates, divisions or subdivisions that are involved in manufacture, distribution and/or sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 50:

All charts or accounts that identify and describe all accounts associated with the development, manufacture, sale, service, distribution and administrative activities associated with the Accused Products.

REQUEST FOR PRODUCTION NO. 51:

All documents that refer or relate to market studies, surveys, analyses, third party industry studies and/or analyst reports related to the Accused Products.

REQUEST FOR PRODUCTION NO. 52:

All drafts, proposals, final copies, drawings, videotapes, audio tapes, electronic documents including web pages for advertising, point-of-sale commercials, or other promotional material for use in the United States for the Accused Products.

REQUEST FOR PRODUCTION NO. 53:

All documents that identify by name, company, address and title, all third parties hired by your or your counsel to investigate this matter or review any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 54:

All documents mentioning or relating to any of the named inventors of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 55:

All documents referring or relating to efforts to design around the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 56:

Documents sufficient to identify whom the Accused Products are sold to and in what quantities.

REQUEST FOR PRODUCTION NO. 57:

Documents sufficient to identify any other litigation and/or settlements relating to network intrusion detection, prevention and analysis products with which you have been involved.

REQUEST FOR PRODUCTION NO. 58:

All documents referring or relating to the accounting policies that are followed with respect to recording sales of the Accused Products.

REQUEST FOR PRODUCTION NO. 59:

All documents referring or relating to the pricing policies for the Accused Products, including and pricing policies between your subsidiaries and/or affiliated entities.

REQUEST FOR PRODUCTION NO. 60:

All documents referring or relating to current inventories of the Accused Products.

REQUEST FOR PRODUCTION NO. 61:

All documents describing or relating to you policies on licensing or cross-licensing technology.

REQUEST FOR PRODUCTION NO. 62:

All documents referring or relating to the installation and use of Accused Products by customers.

REQUEST FOR PRODUCTION NO. 63:

All documents referring to, relating to, or comprising any license entered into or proposed by ISS-GA to obtain rights to make, use or sell the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 64:

All documents referring to, relating to, or comprising any contract or agreement by ISS-GA to take a license to make, use or sell any of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 65:

All documents referring to, relating to or comprising any decision by ISS-GA to issue a license to any person for rights to any of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 66:

All documents referring to, relating to or comprising any request by any person for a license from ISS-GA for any of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 67:

All documents referring or relating to any negotiations for a license under any ISS-GA patent or patent application relating to the Accused Products.

REQUEST FOR PRODUCTION NO. 68:

All documents referring or related to SRI.

REQUEST FOR PRODUCTION NO. 69:

All documents that support your contention that you do not infringe the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 70:

All documents that support your contention that the Patents-in-Suit are invalid.

REQUEST FOR PRODUCTION NO. 71:

All documents that support your contention that the Patents-in-Suit are unenforceable.

REQUEST FOR PRODUCTION NO. 72:

All communications with customers regarding the Accused Products.

REQUEST FOR PRODUCTION NO. 73:

All documents that refer or relate to receiving a request for samples of the Accused Products, and providing samples of the Accused Products, including without limitation requests received through the internet or otherwise and samples provided in response to such requests.

REQUEST FOR PRODUCTION NO. 74:

All documents relating to competition for the Accused Products.

REQUEST FOR PRODUCTION NO. 75:

All documents relating to internal testing of the Accused Products, and any internal ISS-GA Lab used to test the products.

REQUEST FOR PRODUCTION NO. 76:

All documents relating to any third party consultants used by ISS-GA's customers to implement the Accused Products.

REQUEST FOR PRODUCTION NO. 77:

All Annual Reports since 2000.

REQUEST FOR PRODUCTION NO. 78:

The source code for any software or firmware imbedded in or constituting Accused Products.

REQUEST FOR PRODUCTION NO. 79:

All documents relating to communications with third parties regarding SRL.

REQUEST FOR PRODUCTION NO. 80:

All documents referring or relating to the installation, maintenance and/or operation of Accused Products on your customers' computer networks.

///

///

///

///

///

///

///

///


///

REQUEST FOR PRODUCTION NO. 81:

All documents referring or relating to the deployment of Accused Products in the computer networks of your customers.

Dated: June 23, 2005

FISH & RICHARDSON P.C.

By: 
Timothy Devlin (#4241)
John F. Horvath (#4557)
FISH & RICHARDSON P.C.
919 N. Market St., Ste. 1100
P.O. Box 1114
Wilmington, DE 19889-1114
Telephone: (302) 652-5070
Facsimile: (302) 652-0607

Howard G. Pollack (CA Bar No. 162897)
Michael J. Curley (CA Bar No. 230343)
FISH & RICHARDSON P.C.
500 Arguello Street, Suite 500
Redwood City, California 94063
Telephone: (650) 839-5070
Facsimile: (650) 839-5071

Attorneys for Plaintiff
SRI INTERNATIONAL, INC.

CERTIFICATE OF SERVICE

I hereby certify that on the 23rd day of June, 2005, a true and correct copy of the PLAINTIFF SRI INTERNATIONAL, INC.'S FIRST SET OF REQUESTS FOR PRODUCTION TO DEFENDANT INTERNET SECURITY SYSTEMS, INC., A GEORGIA CORPORATION [NOS. 1 – 81] was caused to be served on the attorneys at the following addresses as indicated:

VIA HAND DELIVERY

Richard K. Herrmann
Morris James Hitchens & Williams LLP
111 Delaware Avenue, 10th Floor
Wilmington, DE 19801

Attorney for Defendant
SYMANTEC CORPORATION

VIA HAND DELIVERY

Richard L. Horwitz
Daniel L. Moore
Patterson Anderson & Corroon LLP
Hercules Plaza
133 South Market Street, 6th Floor
P.O. Box 951
Wilmington, DE 19899

Attorneys for Defendant
INTERNET SECURITY SYSTEMS,
INC.

VIA FEDERAL EXPRESS

Paul S. Grewal
Dorcascheer Madrid & Batchelder, LLP
5000 Stevens Creek Boulevard, Suite 400
San Jose, California 95014

Attorney for Defendant
SYMANTEC CORPORATION

VIA FEDERAL EXPRESS

Holmes Hawkins, III
King & Spalding
10 Peachtree Street, N.E.
Atlanta, GA 30303-1763

Attorneys for Defendant
INTERNET SECURITY SYSTEMS,
INC.



John F. Horvath

EXHIBIT E

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

SRI INTERNATIONAL, INC., a California
Corporation,

Plaintiff,

v.

INTERNET SECURITY SYSTEMS, INC.,
a Delaware corporation, and INTERNET
SECURITY SYSTEMS, INC., a Georgia
corporation, and SYMANTEC
CORPORATION, a Delaware corporation,

Defendants.

Case No. 04-1199-SLR

**PLAINTIFF SRI INTERNATIONAL, INC.'S FIRST SET OF REQUESTS FOR
PRODUCTION TO DEFENDANT SYMANTEC CORPORATION [NOS. 1 – 88]**

Pursuant to Rules 26 and 34 of the Federal Rules of Civil Procedure, Plaintiff SRI International, Inc. hereby requests Defendant Symantec Corporation, to respond to these Requests by producing for inspection and copying the Documents and Things requested below at the offices of Fish & Richardson P.C., 500 Arguello Street, Suite 500, Redwood City, California 94063, within thirty (30) days of service of these Requests, or at such other times and places as counsel for the parties may agree.

INSTRUCTIONS

A. When producing a Document, please produce the Document as it is kept in the ordinary course of business, or indicate the paragraph of these Requests to which that Document is responsive.

B. Electronic records and computerized information must be produced in an intelligible format or together with a description of the system from which it was derived sufficient to permit rendering the materials intelligible.

C. In producing Documents, furnish all Documents known or available to Symantec regardless of whether such Documents are possessed directly by Symantec, or by any parent, subsidiary, or affiliated corporation, or any of Symantec's officers, directors, employees, agents, representatives, or attorneys, as well as any other Documents in Symantec's custody or control.

D. File folders with tabs or labels identifying Documents called for by these Requests must be produced intact with such Documents.

E. Selection of Documents from the files and other sources and the numbering of such Documents shall be performed in such a manner as to insure that the source of each Document may be determined, if necessary.

F. Documents attached to each other must not be separated.

G. The term "all Documents" means any and all Documents that might reasonably be located through a search of all locations reasonably likely to contain Documents called for by these Requests.

H. Should any Document be withheld based on some limitation of discovery (Including a claim of privilege), please supply the following information:

1. The identity(ies) of each Document's author(s), writer(s), sender(s), or initiator(s);
2. The identity(ies) of each Document's recipient(s), addressee(s), or party(ies) for whom it was intended;
3. The date of creation or transmittal indicated on each Document, or an estimate of that date, indicated as such, if no date appears on the Document;
4. The general subject matter as described on each Document, or, if no such description appears, then some other description sufficient to identify the Document; and
5. The claimed ground(s) for limitation of discovery (e.g., "attorney-client privilege" or "attorney work product doctrine").

I. The written answer to each individual request for production must repeat verbatim, immediately before each answer, the text of the individual request for production being answered.

DEFINITIONS

As used in these Requests, the following terms have the meanings indicated:

A. "SRP" means SRI International, Inc., including its officers, directors, employees, agents, and attorneys.

B. "Symantec," "Defendant," "you," or "your" means Symantec Corporation, including its past and present officers, directors, employees, consultants, agents, and attorneys and others acting or purporting to act on its behalf, and including their predecessors, subsidiaries, parents, and affiliates.

C. The phrase the "'615 patent" means U.S. Patent No. 6,711,615.

D. The phrase the "'203 patent" means U.S. Patent No. 6,484,203.

E. The phrase the "'338 patent" means U.S. Patent No. 6,321,338.

F. The phrase the "'212 patent" means U.S. Patent No. 6,708,212.

G. The phrase the "'874 patent" means U.S. Patent No. 6,704,874.

H. The phrase "Patents-in-Suit" means the '615, '203, '338, '212, and '874 Patents.

I. The word "Document" is used herein in its broadest sense to include everything that is contemplated by Rule 26 and Rule 34 of the Federal Rules of Civil Procedure, Including Documents stored in hard copy or electronic form. Electronic Documents Include electronic mail, computer source code, object code, and microcode, and Documents stored on any media accessible by electronic means. A comment or notation appearing on any "Document" and not a part of the original text is to be considered a separate "Document."

J. "Thing" means any tangible object other than a Document.

K. "Person" or "Persons" include not only natural individuals, but also, without limitation, firms, partnerships, associations, corporations, and other legal entities, and divisions, departments, or other units thereof.

L. These "Requests" shall mean the instant document, i.e., Plaintiff SRI International, Inc.'s First Set of Requests for Production of Documents and Things to Defendant Symantec Corporation [Nos. 1 – 88].

M. "Including" shall mean "including but not limited to."

N. The terms "and" and "or" shall be construed conjunctively or disjunctively, whichever makes the individual request more inclusive.

O. The singular and masculine form of any noun or pronoun shall embrace and be read and applied as embracing the plural, the feminine, and the neuter, except where circumstances clearly make it inappropriate.

P. The terms "refer," "referring," "relate," or "relating" as used herein include, but are not limited to the following meanings: bearing upon, concerning, constituting, discussing, describing, evidencing, identifying, concerning, mentioning, in connection with, pertaining to, respecting, regarding, responding to, or in any way factually or logically relevant to the matter described in the request.

Q. The term "Accused Products" shall mean all software products relating to computer network intrusion detection, prevention or analysis developed, made, used, sold or offered for sale in, or imported into the United States at any time from 2001 to the present, including but not limited to ManHunt, Symantec Security Gateway, Symantec EventManager for Intrusion Protection, and Symantec Security Management System.

DOCUMENTS AND THINGS TO BE PRODUCED

REQUEST FOR PRODUCTION NO. 1:

All documents referring or relating in any way to any of your document, file, or record retention or destruction policies, including without limitation, any such policies referring or relating to electronic data, including source code, from 2001 to the present.

REQUEST FOR PRODUCTION NO. 2:

Management and/or organization charts setting forth the entities and/or the names and/or titles of individuals involved in the decision to develop, development, manufacture, marketing, sale and distribution of the Accused Products.

REQUEST FOR PRODUCTION NO. 3:

All brochures, advertisements, press releases, price lists, catalogs and other marketing and promotional materials relating to the Accused Products.

REQUEST FOR PRODUCTION NO. 4:

All documents containing, describing or relating to any communication between you and any third-party supplier of pieces of software, including source code, or other technology incorporated in the Accused Products.

REQUEST FOR PRODUCTION NO. 5:

All documents that describe, illustrate, or depict names and/or functions of subsidiaries, departments or affiliated entities related to you that were involved in any matter in the decision to develop, conception, design, manufacture, testing, marketing or sale of any of the Accused Products.

REQUEST FOR PRODUCTION NO. 6:

All documents that refer or relate to the research, development, and testing of the Accused Products, including but not limited to drawings, prototypes, notes, notebooks, workbooks, project reports, correspondence, memoranda, test results, schematics, flow charts and invention disclosures.

REQUEST FOR PRODUCTION NO. 7:

All operation manuals, user manuals, installation guides, dealers guides, and service manuals for the Accused Products.

REQUEST FOR PRODUCTION NO. 8:

All documents relating to the decision to develop and market the Accused Products, including all market studies, communication with customers, and internal correspondence.

REQUEST FOR PRODUCTION NO. 9:

All documents that refer or relate to the design of the Accused Products, including but not limited to specifications, data sheets, drawings, diagrams, schematics, manuals, notes, notebooks, workbooks, correspondence, and memoranda.

REQUEST FOR PRODUCTION NO. 10:

All documents sufficient to show the structure, function, or operation of the Accused Products, including but not limited to drawings, , block diagrams and operating instructions.

REQUEST FOR PRODUCTION NO. 11:

All documents that refer or relate to the research, development, and testing of the method and/or process for using any Accused Product or products incorporating the same, including but not limited to notes, notebooks, workbooks, correspondence, memoranda, and test results.

REQUEST FOR PRODUCTION NO. 12:

All documents that refer or relate to the design of the method and/or process for using any Accused Product or products incorporating the same, including but not limited to instructions for use, drawings, , diagrams, manuals, notes, notebooks, workbooks, correspondence, and memoranda.

REQUEST FOR PRODUCTION NO. 13:

All documents that refer or relate to the manufacture, , and production of the Accused Products, from 2001 to the present, including but not limited to the location of all manufacturing, and production facilities and the quantity of each Accused Product manufactured or produced at each such facility.

REQUEST FOR PRODUCTION NO. 14:

All documents, including but not limited to promotional materials, press releases, trade journal articles, advertisements, catalogs, documents prepared for trade shows and meetings, package inserts, technical data sheets, specifications, price lists, sales presentations, and sales and marketing forecasts and projections, that refer or relate to advertising and marketing of the Accused Products from 2001 to the present.

REQUEST FOR PRODUCTION NO. 15:

All documents that refer or relate to the sale within the United States of the Accused Products, from 2001 to the present, including but not limited to documents sufficient to show Symantec's actual and expected gross revenues, net profits, gross revenues, sales, sales prices, and returns of the Accused Products.

REQUEST FOR PRODUCTION NO. 16:

All documents that have been or are included with the sale of the Accused Products or products containing same, including without limitation instruction manuals, instructions for use, and users' guides. .

REQUEST FOR PRODUCTION NO. 17:

All documents referring or relating to the distribution network maintained by Symantec for purposes of selling, distributing, or licensing the Accused Products and products containing same.

REQUEST FOR PRODUCTION NO. 18:

All documents that refer or relate to the importation into the United States of the Accused Products and any product incorporating an Accused Product, from 2001 to the present, including but not limited to the quantity of Accused Products (specifically identified by product) imported into the United States, the dates of such importation, and the ports of entry.

REQUEST FOR PRODUCTION NO. 19:

All documents that refer or relate to any United States inventory of the Accused Products, from 2001 to the present, including but not limited to the location of inventory, the amount of inventory, and the owner and/or holder of each such inventory.

REQUEST FOR PRODUCTION NO. 20:

All documents sufficient to show or relate to any step in the method and/or process of using the Accused Products and products incorporating the Accused Products, including but not limited to drawings, and block diagrams.

REQUEST FOR PRODUCTION NO. 21:

Documents sufficient to identify each and every Accused Product that has been imported into the United States, sold for importation into the United States, or sold or offered for sale after importation into the United States, from 2001 to the present.

REQUEST FOR PRODUCTION NO. 22:

Documents sufficient to identify each and every Accused Product that is being planned, that has been designed, and/or that is in development that is capable of being used in a product, circuit, or device that has been imported into the United States, sold for importation into the United States, or sold or offered for sale after importation into the United States.

REQUEST FOR PRODUCTION NO. 23:

Five samples of each version of each Accused Product, and all literature or documentation distributed with each Accused Product, any additional products that are required for the intended uses (both inside the United States and outside the United States) of each Accused Product, and all instruction or other guidance information associated with the sale, demonstration or use of each Accused Product.

REQUEST FOR PRODUCTION NO. 24:

All documents that refer or relate to instructions for using the Accused Products, including prior versions of instructions for use and instructions for use provided with the

sale, demonstration, marketing, or training for use of the Accused Products, both inside the United States and outside the United States.

REQUEST FOR PRODUCTION NO. 25:

All documents discussing or making reference to the quality, value, acceptability, workability, performance or benefits of any Accused Product, including communications with anyone who has used a Accused Product integrated circuit product praising, criticizing, or otherwise commenting on the Accused Products.

REQUEST FOR PRODUCTION NO. 26:

All documents that refer or relate to any failures, problems with and/or complaints about any of the Accused Products and products containing same or the method and/or process of using any Accused Product and products containing same.

REQUEST FOR PRODUCTION NO. 27:

All documents that refer or relate to uses, tests, or evaluations of the Accused Products and products containing same or the method and/or process of using any Accused Product and products containing same performed or conducted by Symantec, Symantec's customers, or any other third party that Symantec is aware of.

REQUEST FOR PRODUCTION NO. 28:

All documents that refer or relate to training any customer in the use of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 29:

All documents evidencing contributions made by any of your employees to the conception, design, or development of the Accused Products.

REQUEST FOR PRODUCTION NO. 30:

All United States and foreign patents and patent applications, including all documents referring or relating to such patents and patent applications, owned or controlled by you that relate to the Accused Products.

REQUEST FOR PRODUCTION NO. 31:

All documents constituting or relating to licenses or indemnification agreements between you and any other party relating to the Accused Products, or comparable products.

REQUEST FOR PRODUCTION NO. 32:

All articles, speeches, presentations or interviews that have been written or given by your employees, officers, directors or other representatives that refer or relate to the Accused Products.

REQUEST FOR PRODUCTION NO. 33:

All documents referring or relating to the Patents-in-Suit, including, but not limited to, documents that refer or relate to when and how you first became aware of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 34:

All documents found or identified during any enforceability searches, infringement analyses, prior art searches or studies related to the Patents-in-Suit, including any copies of patents, publications or other prior art.

REQUEST FOR PRODUCTION NO. 35:

All documents referring or relating to or constituting any patent or publication that you contend invalidates any of the claims of any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 36:

All documents, including but not limited to opinion letters, memoranda, or other documents that refer to or relate to the validity, infringement and/or enforceability of any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 37:

All internal communications among and between your personnel that reflect, refer to, or concern any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 38:

All documents, including without limitation articles and internal technical memoranda authored by you that relate to the subject matter described, disclosed or claimed in any of the patents-in-suit.

REQUEST FOR PRODUCTION NO. 39:

All communications or opinions of your officers, directors and/or employees regarding the infringement, validity or enforceability of any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 40:

All documents referring or relating to any consideration given by Symantec to any possible liability for patent infringement should Symantec commence or continue to make, use, distribute or sell the Accused Products or products containing same, including without limitation any internal notes or memoranda.

REQUEST FOR PRODUCTION NO. 41:

All documents that set forth on a monthly, quarterly or annual basis from 2001 to the present, the number and nature of units developed, used, sold and/or distributed for each of the Accused Products.

REQUEST FOR PRODUCTION NO. 42:

All documents that set forth, on a monthly, quarterly or annual basis from 2001 to present, the amount of sales in U.S. dollars or local currency (indicate currency type if not U.S. dollars) for each of the Accused Products.

REQUEST FOR PRODUCTION NO. 43:

All projections, forecasts, market share reports, marketing acceptance documents business plans, strategic plans, fiscal plans, marketing plans, or sales plans that refer or relate to the sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 44:

All documents that, on a monthly, quarterly or annual basis, refer or relate to unit sales, gross selling price, net selling price, or amounts deducted from gross selling price to arrive at net selling price, for each of the Accused Products.

REQUEST FOR PRODUCTION NO. 45:

All profit and loss statements, printouts, statements or other documents that set forth or include gross profit figures for the Accused Products on a monthly, quarterly or annual basis from 2001 to present.

REQUEST FOR PRODUCTION NO. 46:

All annual, semi-annual, quarterly, monthly or other documents that set forth or include your gross expenses incurred in the manufacture, distribution and sale of the Accused Products, including but not limited to, direct labor costs, direct manufacturing costs, selling costs, variable overhead costs, and all other costs associated with the manufacture, distribution and sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 47:

Documents that evidence the date of the first sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 48:

All documents that identify any of your divisions, affiliates, subsidiaries, or suppliers that have developed or manufactured any of the Accused Products.

REQUEST FOR PRODUCTION NO. 49:

Financial statements, including profit and loss statements, income statements, balance sheets, statements of cash flow, statements of retained earnings and notes thereto for any of your affiliates, divisions or subdivisions that are involved in manufacture, distribution and/or sale of the Accused Products.

REQUEST FOR PRODUCTION NO. 50:

All charts or accounts that identify and describe all accounts associated with the development, manufacture, sale, service, distribution and administrative activities associated with the Accused Products.

REQUEST FOR PRODUCTION NO. 51:

All documents that refer or relate to market studies, surveys, analyses, third party industry studies and/or analyst reports related to the Accused Products.

REQUEST FOR PRODUCTION NO. 52:

All drafts, proposals, final copies, drawings, videotapes, audio tapes, electronic documents including web pages for advertising, point-of-sale commercials, or other promotional material for use in the United States for the Accused Products.

REQUEST FOR PRODUCTION NO. 53:

All documents that identify by name, company, address and title, all third parties hired by your or your counsel to investigate this matter or review any of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 54:

All documents mentioning or relating to any of the named inventors of the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 55:

All documents referring or relating to efforts to design around the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 56:

Documents sufficient to identify whom the Accused Products are sold to and in what quantities.

REQUEST FOR PRODUCTION NO. 57:

Documents sufficient to identify any other litigation and/or settlements relating to network intrusion detection, prevention and analysis products with which you have been involved.

REQUEST FOR PRODUCTION NO. 58:

All documents referring or relating to the accounting policies that are followed with respect to recording sales of the Accused Products.

REQUEST FOR PRODUCTION NO. 59:

All documents referring or relating to the pricing policies for the Accused Products, including and pricing policies between your subsidiaries and/or affiliated entities.

REQUEST FOR PRODUCTION NO. 60:

All documents referring or relating to current inventories of the Accused Products.

REQUEST FOR PRODUCTION NO. 61:

All documents describing or relating to you policies on licensing or cross-licensing technology.

REQUEST FOR PRODUCTION NO. 62:

All documents referring or relating to the installation and use of Accused Products by customers.

REQUEST FOR PRODUCTION NO. 63:

All documents referring to, relating to, or comprising any license entered into or proposed by Symantec to obtain rights to make, use or sell the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 64:

All documents referring to, relating to, or comprising any contract or agreement by Symantec to take a license to make, use or sell any of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 65:

All documents referring to, relating to or comprising any decision by Symantec to issue a license to any person for rights to any of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 66:

All documents referring to, relating to or comprising any request by any person for a license from Symantec for any of the Accused Products or products containing same.

REQUEST FOR PRODUCTION NO. 67:

All documents referring or relating to any negotiations for a license under any Symantec patent or patent application relating to the Accused Products.

REQUEST FOR PRODUCTION NO. 68:

All documents referring or related to SRI.

REQUEST FOR PRODUCTION NO. 69:

All documents that support your contention that you do not infringe the Patents-in-Suit.

REQUEST FOR PRODUCTION NO. 70:

All documents that support your contention that the Patents-in-Suit are invalid.

REQUEST FOR PRODUCTION NO. 71:

All documents that support your contention that the Patents-in-Suit are unenforceable.

REQUEST FOR PRODUCTION NO. 72:

All communications with customers regarding the Accused Products.

REQUEST FOR PRODUCTION NO. 73:

All documents that refer or relate to receiving a request for samples of the Accused Products, and providing samples of the Accused Products, including without

limitation requests received through the internet or otherwise and samples provided in response to such requests.

REQUEST FOR PRODUCTION NO. 74:

All documents relating to competition for the Accused Products.

REQUEST FOR PRODUCTION NO. 75:

All documents relating to internal testing of the Accused Products, and any internal Symantec Lab used to test the products.

REQUEST FOR PRODUCTION NO. 76:

All documents relating to any third party consultants used by Symantec's customers to implement the Accused Products.

REQUEST FOR PRODUCTION NO. 77:

All Annual Reports since 2000.

REQUEST FOR PRODUCTION NO. 78:

The source code for any software or firmware imbedded in or constituting Accused Products.

REQUEST FOR PRODUCTION NO. 79:

All documents relating to communications with third parties regarding SRI.

REQUEST FOR PRODUCTION NO. 80:

All documents referring or relating to the installation, maintenance and/or operation of Accused Products on your customers' computer networks.

///

///

///

///

///

///

///

REQUEST FOR PRODUCTION NO. 81:

All documents referring or relating to the deployment of Accused Products in the computer networks of your customers.

Dated: June 23, 2005

FISH & RICHARDSON P.C.

By: 

Timothy Devlin (#4241)
John F. Horvath (#4557)
FISH & RICHARDSON P.C.
919 N. Market St., Ste. 1100
P.O. Box 1114
Wilmington, DE 19889-1114
Telephone: (302) 652-5070
Facsimile: (302) 652-0607

Howard G. Pollack (CA Bar No. 162897)
Michael J. Curley (CA Bar No. 230343)
FISH & RICHARDSON P.C.
500 Arguello Street, Suite 500
Redwood City, California 94063
Telephone: (650) 839-5070
Facsimile: (650) 839-5071

Attorneys for Plaintiff
SRI INTERNATIONAL, INC.

RFPs to Symantec (3).doc

CERTIFICATE OF SERVICE

I hereby certify that on the 23rd day of June, 2005, a true and correct copy of the
**PLAINTIFF SRI INTERNATIONAL, INC.'S FIRST SET OF REQUESTS FOR
PRODUCTION TO DEFENDANT SYMANTEC CORPORATION [NOS. 1 – 88]**
was caused to be served on the attorneys of record at the following addresses as
indicated:

VIA HAND DELIVERY

Richard K. Herrmann
Morris, James Hitchens & Williams LLP
222 Delaware Avenue, 10th Floor
Wilmington, DE 19801

Attorney for Defendant
SYMANTEC CORPORATION

VIA HAND DELIVERY

Richard L. Horwitz
David E. Moore
Potter Anderson & Corroon LLP
Hercules Plaza
1313 North Market Street, 6th Floor
P.O. Box 951
Wilmington, DE 19899

Attorneys for Defendant
INTERNET SECURITY SYSTEMS,
INC.

VIA FEDERAL EXPRESS

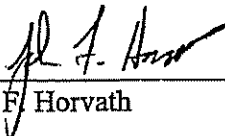
Paul S. Grewal
Day Casebeer Madrid & Batchelder, LLP
20300 Stevens Creek Boulevard, Suite 400
Cupertino, California 95014

Attorney for Defendant
SYMANTEC CORPORATION

VIA FEDERAL EXPRESS

Holmes Hawkins, III
King & Spalding
191 Peachtree Street, N.E.
Atlanta, GA 30303-1763

Attorneys for Defendant
INTERNET SECURITY SYSTEMS,
INC.



John F. Horvath